



US007043078B2

(12) **United States Patent**  
**Guleryuz**

(10) **Patent No.:** **US 7,043,078 B2**  
(45) **Date of Patent:** **\*May 9, 2006**

(54) **METHOD AND APPARATUS FOR SEGMENTATION OF COMPOUND DOCUMENTS HAVING LOW RESOLUTION HALFTONES**

(75) Inventor: **Onur G. Guleryuz**, Sunnyvale, CA (US)

(73) Assignee: **Seiko Epson Corporation**, Tokyo (JP)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 157 days.

This patent is subject to a terminal disclaimer.

4,288,821 A *	9/1981	Lavallee et al. ....	358/3.07
4,893,188 A	1/1990	Murakami et al. ....	358/2.1
5,014,124 A	5/1991	Fujisawa .....	358/530
5,131,059 A	7/1992	Kobayashi et al. ....	358/2.1
5,327,262 A *	7/1994	Williams .....	358/462
5,392,365 A	2/1995	Steinkirchner .....	382/260
5,617,216 A	4/1997	Wada .....	358/2.1
5,949,964 A *	9/1999	Clouthier et al. ....	358/3.06
6,198,850 B1	3/2001	Banton .....	382/239
6,229,923 B1 *	5/2001	Williams et al. ....	382/224
6,347,153 B1 *	2/2002	Triplett et al. ....	382/224
6,385,344 B1	5/2002	Irie et al. ....	382/237
6,389,167 B1	5/2002	Wetchler et al. ....	382/199
6,437,878 B1	8/2002	Hansen et al. ....	358/2.1
2001/0015826 A1	8/2001	Kurose	
2002/0057459 A1	5/2002	Nagarajan et al.	
2002/0067509 A1	6/2002	Roylance	

(21) Appl. No.: **10/729,663**

(22) Filed: **Dec. 5, 2003**

(65) **Prior Publication Data**

US 2004/0175037 A1 Sep. 9, 2004

**Related U.S. Application Data**

(63) Continuation-in-part of application No. 10/625,107, filed on Jul. 23, 2003.

(60) Provisional application No. 60/453,077, filed on Mar. 6, 2003.

(51) **Int. Cl.**

**G06K 9/34** (2006.01)

(52) **U.S. Cl.** ..... **382/180; 382/254; 382/264; 358/3.01; 358/3.27; 358/463**

(58) **Field of Classification Search** ..... 358/1.9, 358/534, 465, 448, 443, 1.2, 3.01, 3.06, 3.07, 358/447, 466, 2.1, 3.27, 463; 382/180, 254, 382/270, 176, 173, 299, 298, 276, 266, 275, 382/264; 328/180, 254

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS


\* cited by examiner

Primary Examiner—Sheela Chawan

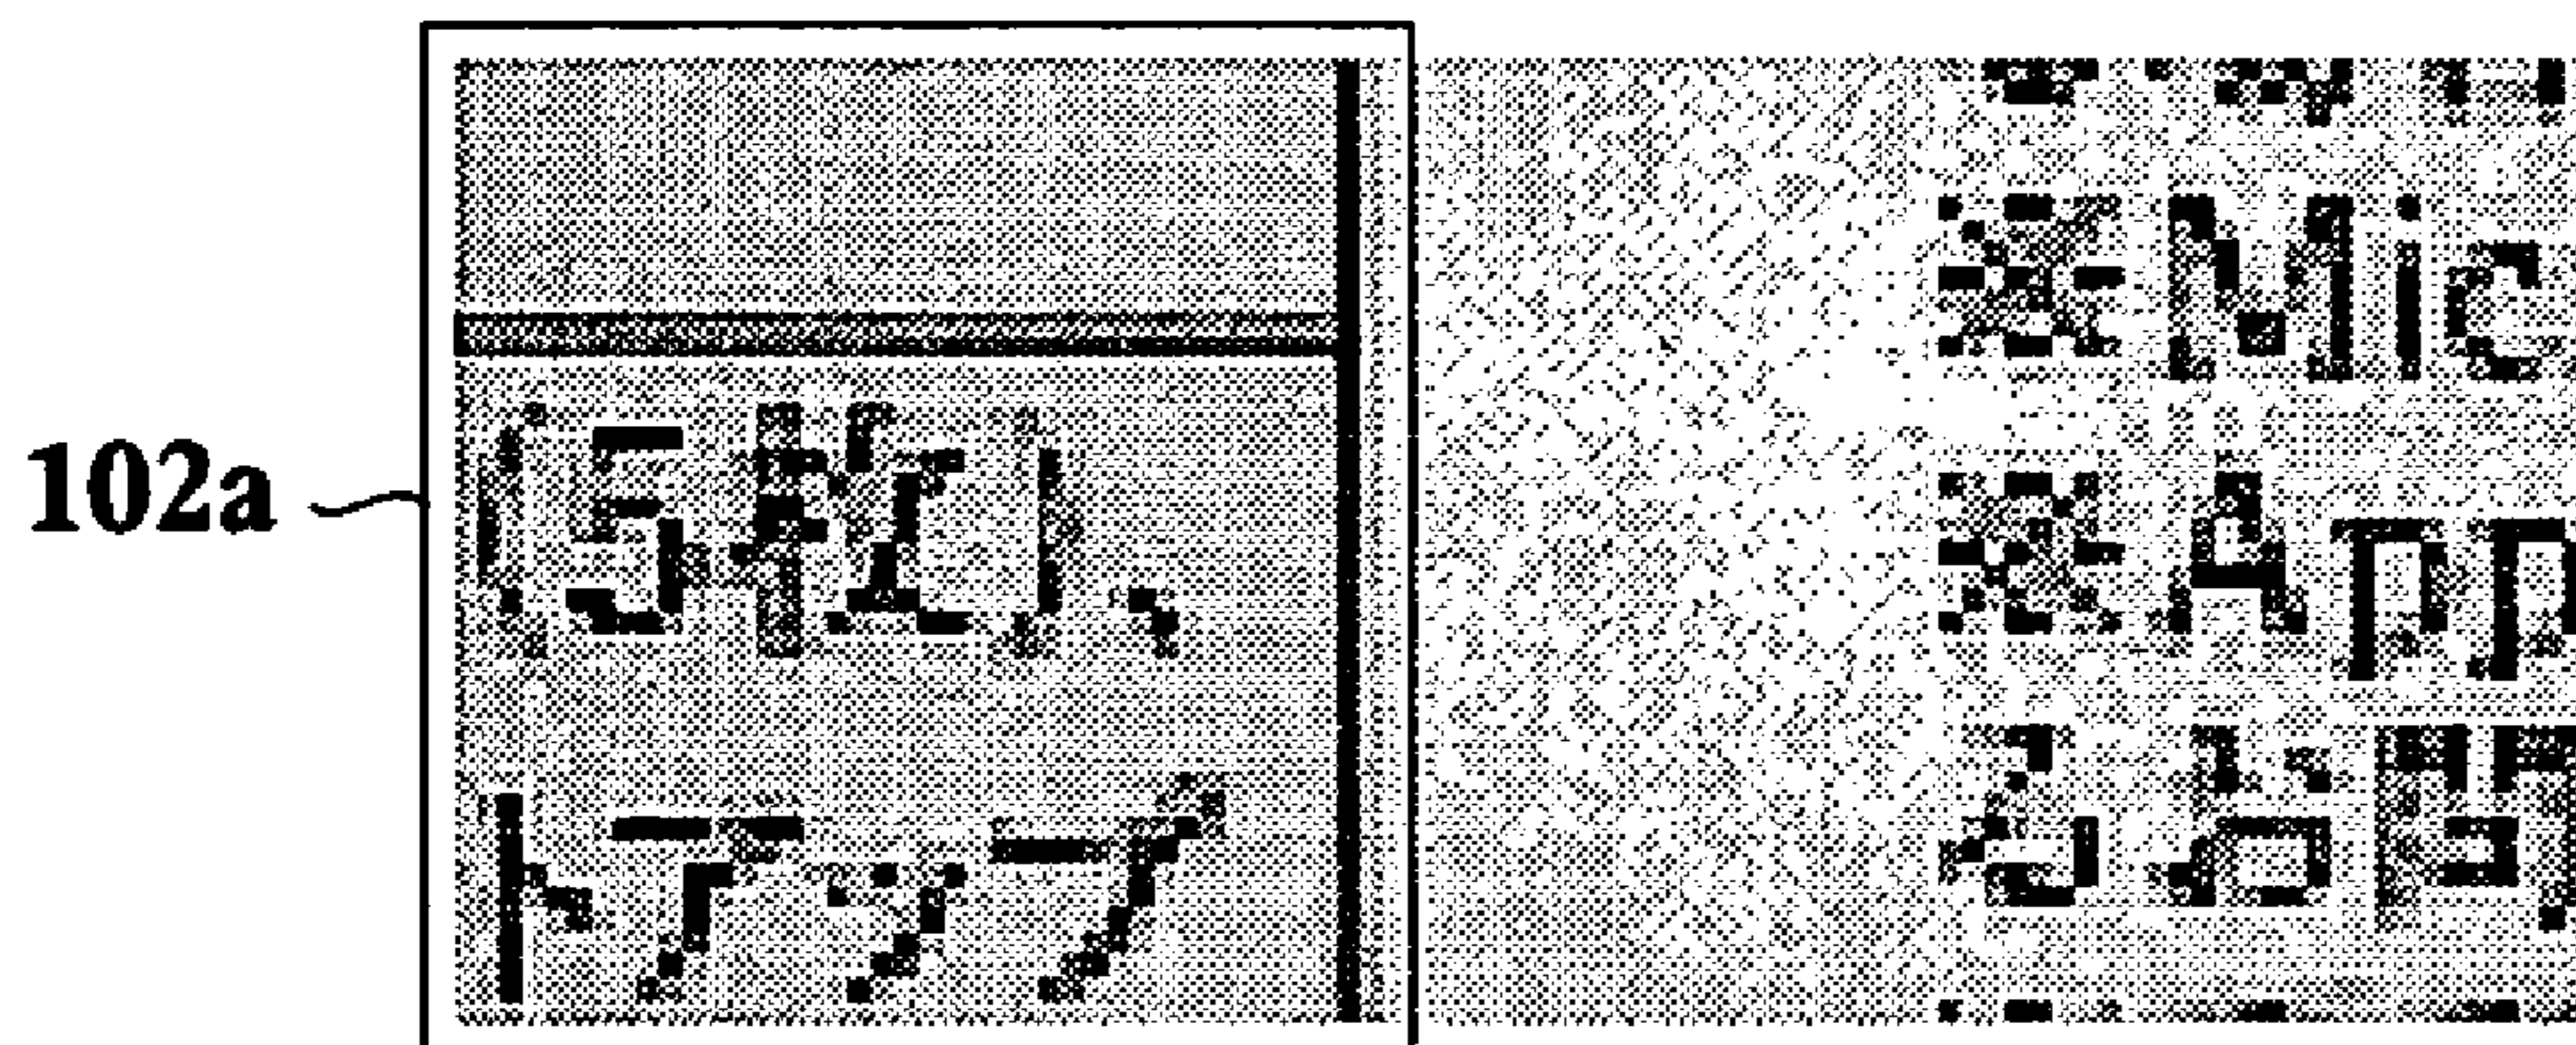
(57) **ABSTRACT**

A method for segmenting a compound document for enhancement during replication of the compound document is provided. The method initiates with filtering data representing a portion of the compound document through a filter associated with a first resolution. Then, the data representing the portion of the compound document is filtered through a filter associated with a second resolution. Next, edges are detected on both, an output of the filter associated with the first resolution and an output of the filter associated with the second resolution. Then, data representing detected edges from both outputs are combined. Next, it is determined whether a pixel corresponding to the data representing the detected edges is over a halftone region. Methods for labeling digital image data and labeling and enhancing documents defined through digital data, as well as associated computer readable medium embodiments are provided. An image replication device and an integrated circuit configured to segment and enhance image data associated with a compound document are also provided.

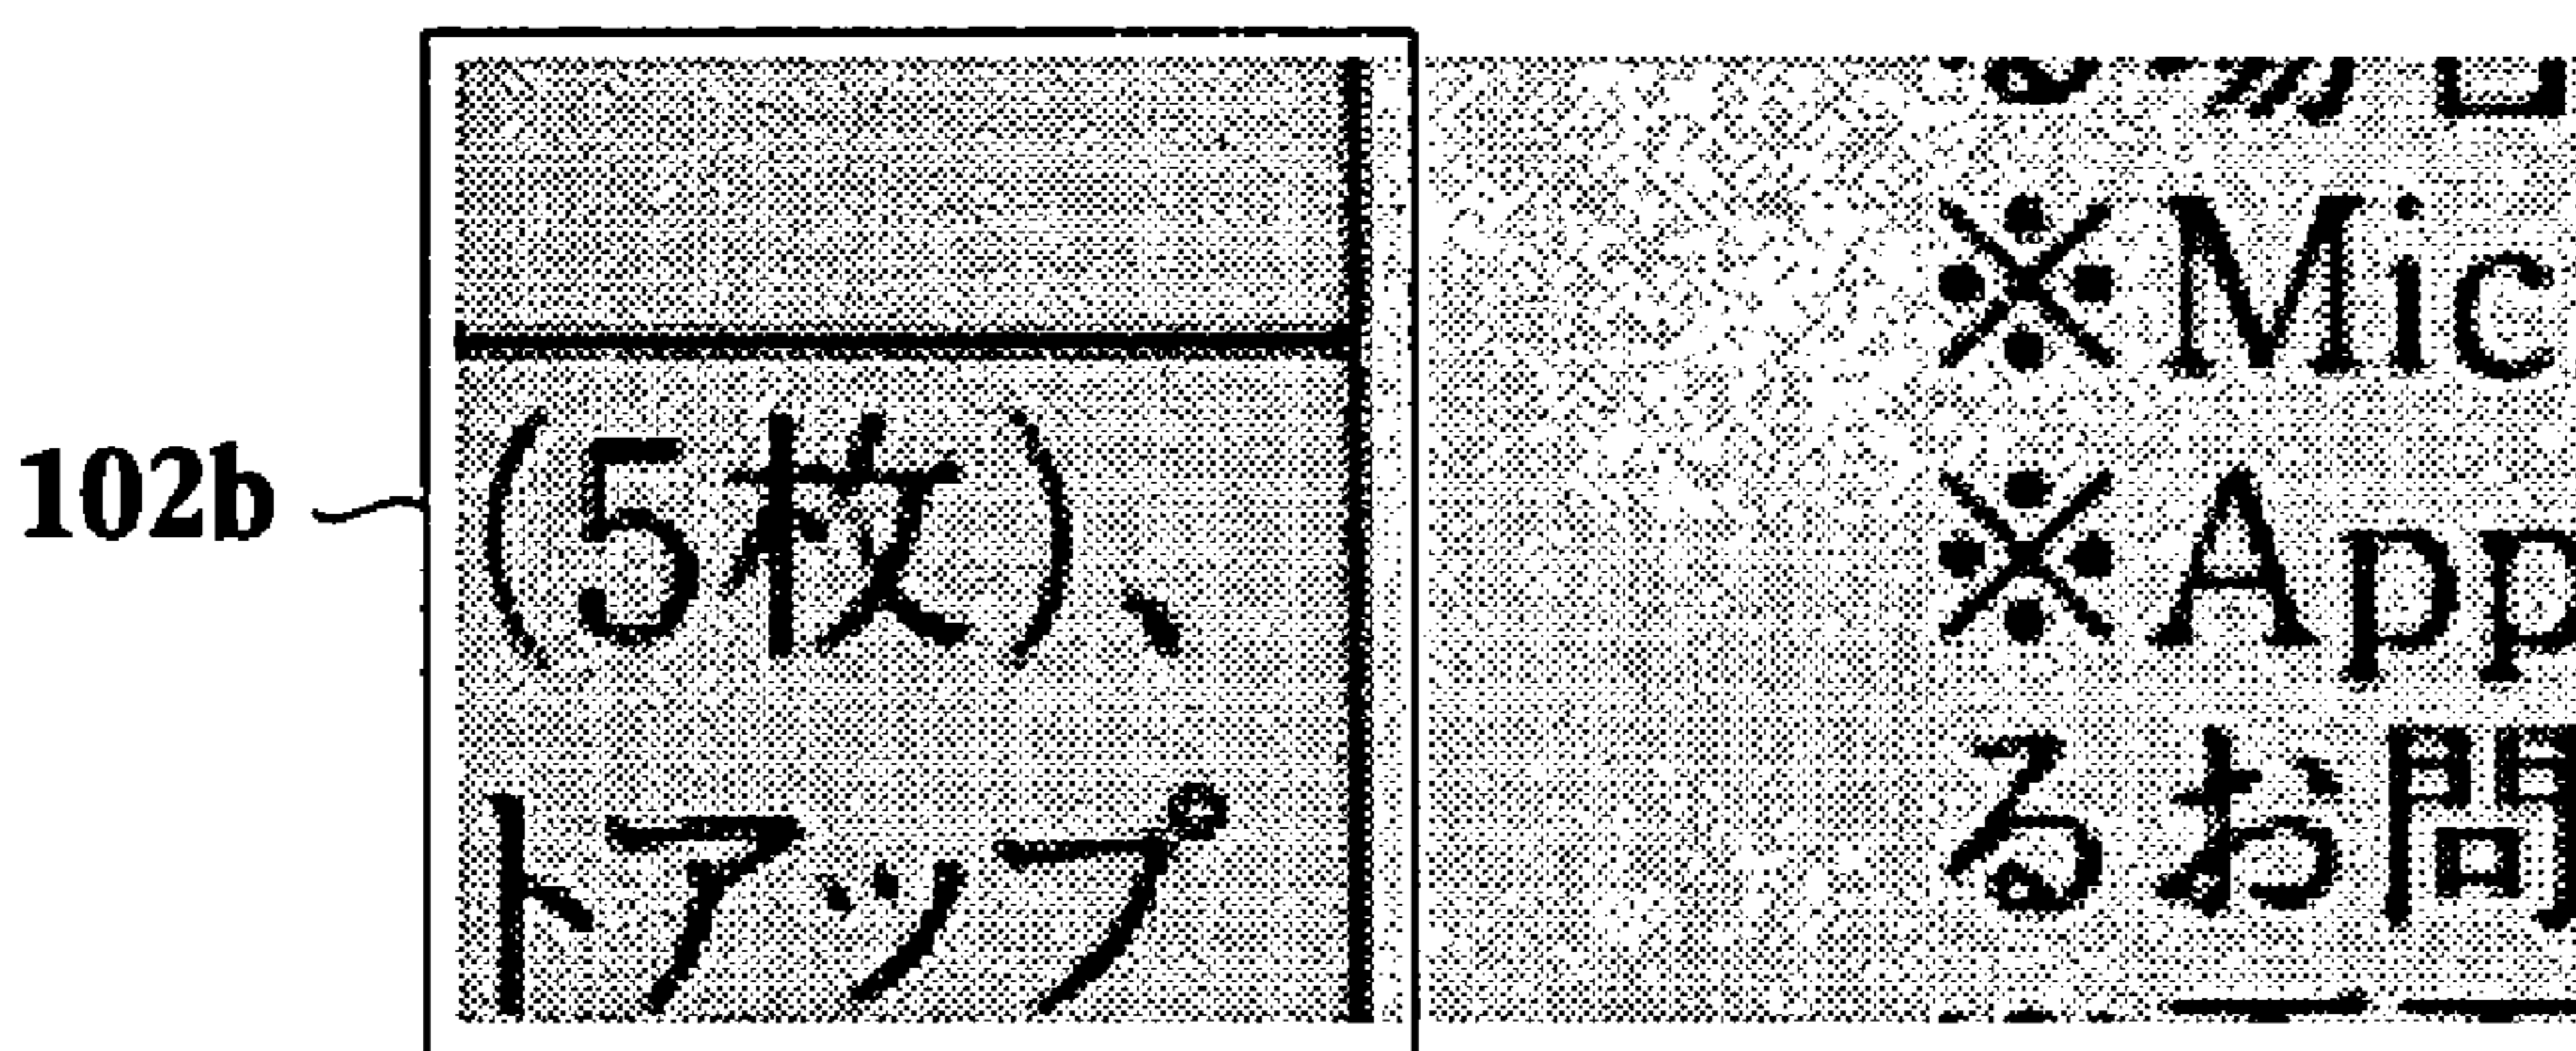
**29 Claims, 32 Drawing Sheets**



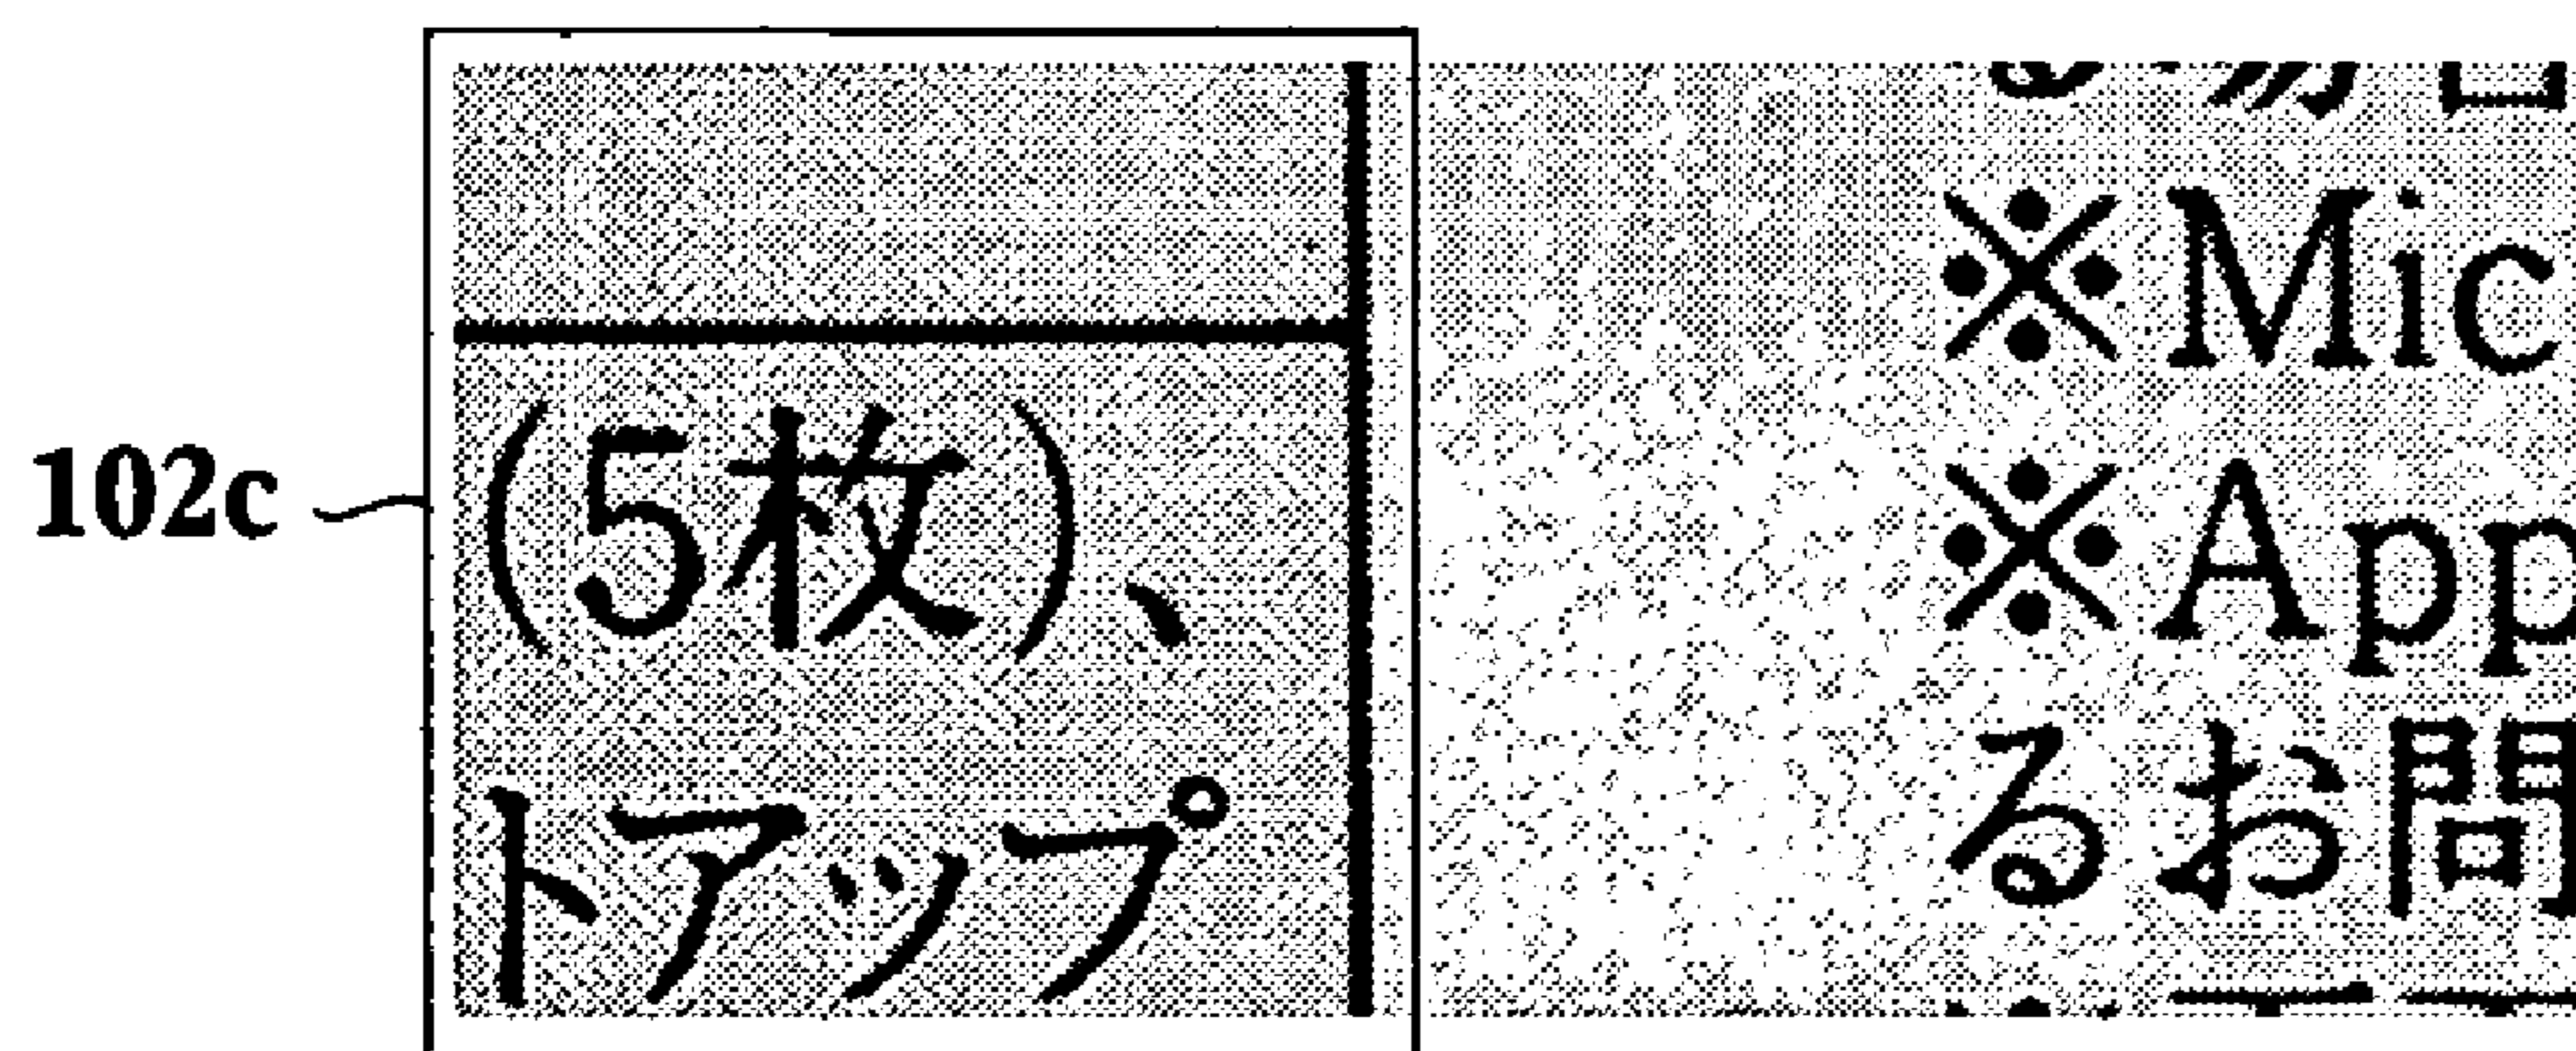




**Fig. 1A**



**Fig. 1B**



**Fig. 1C**

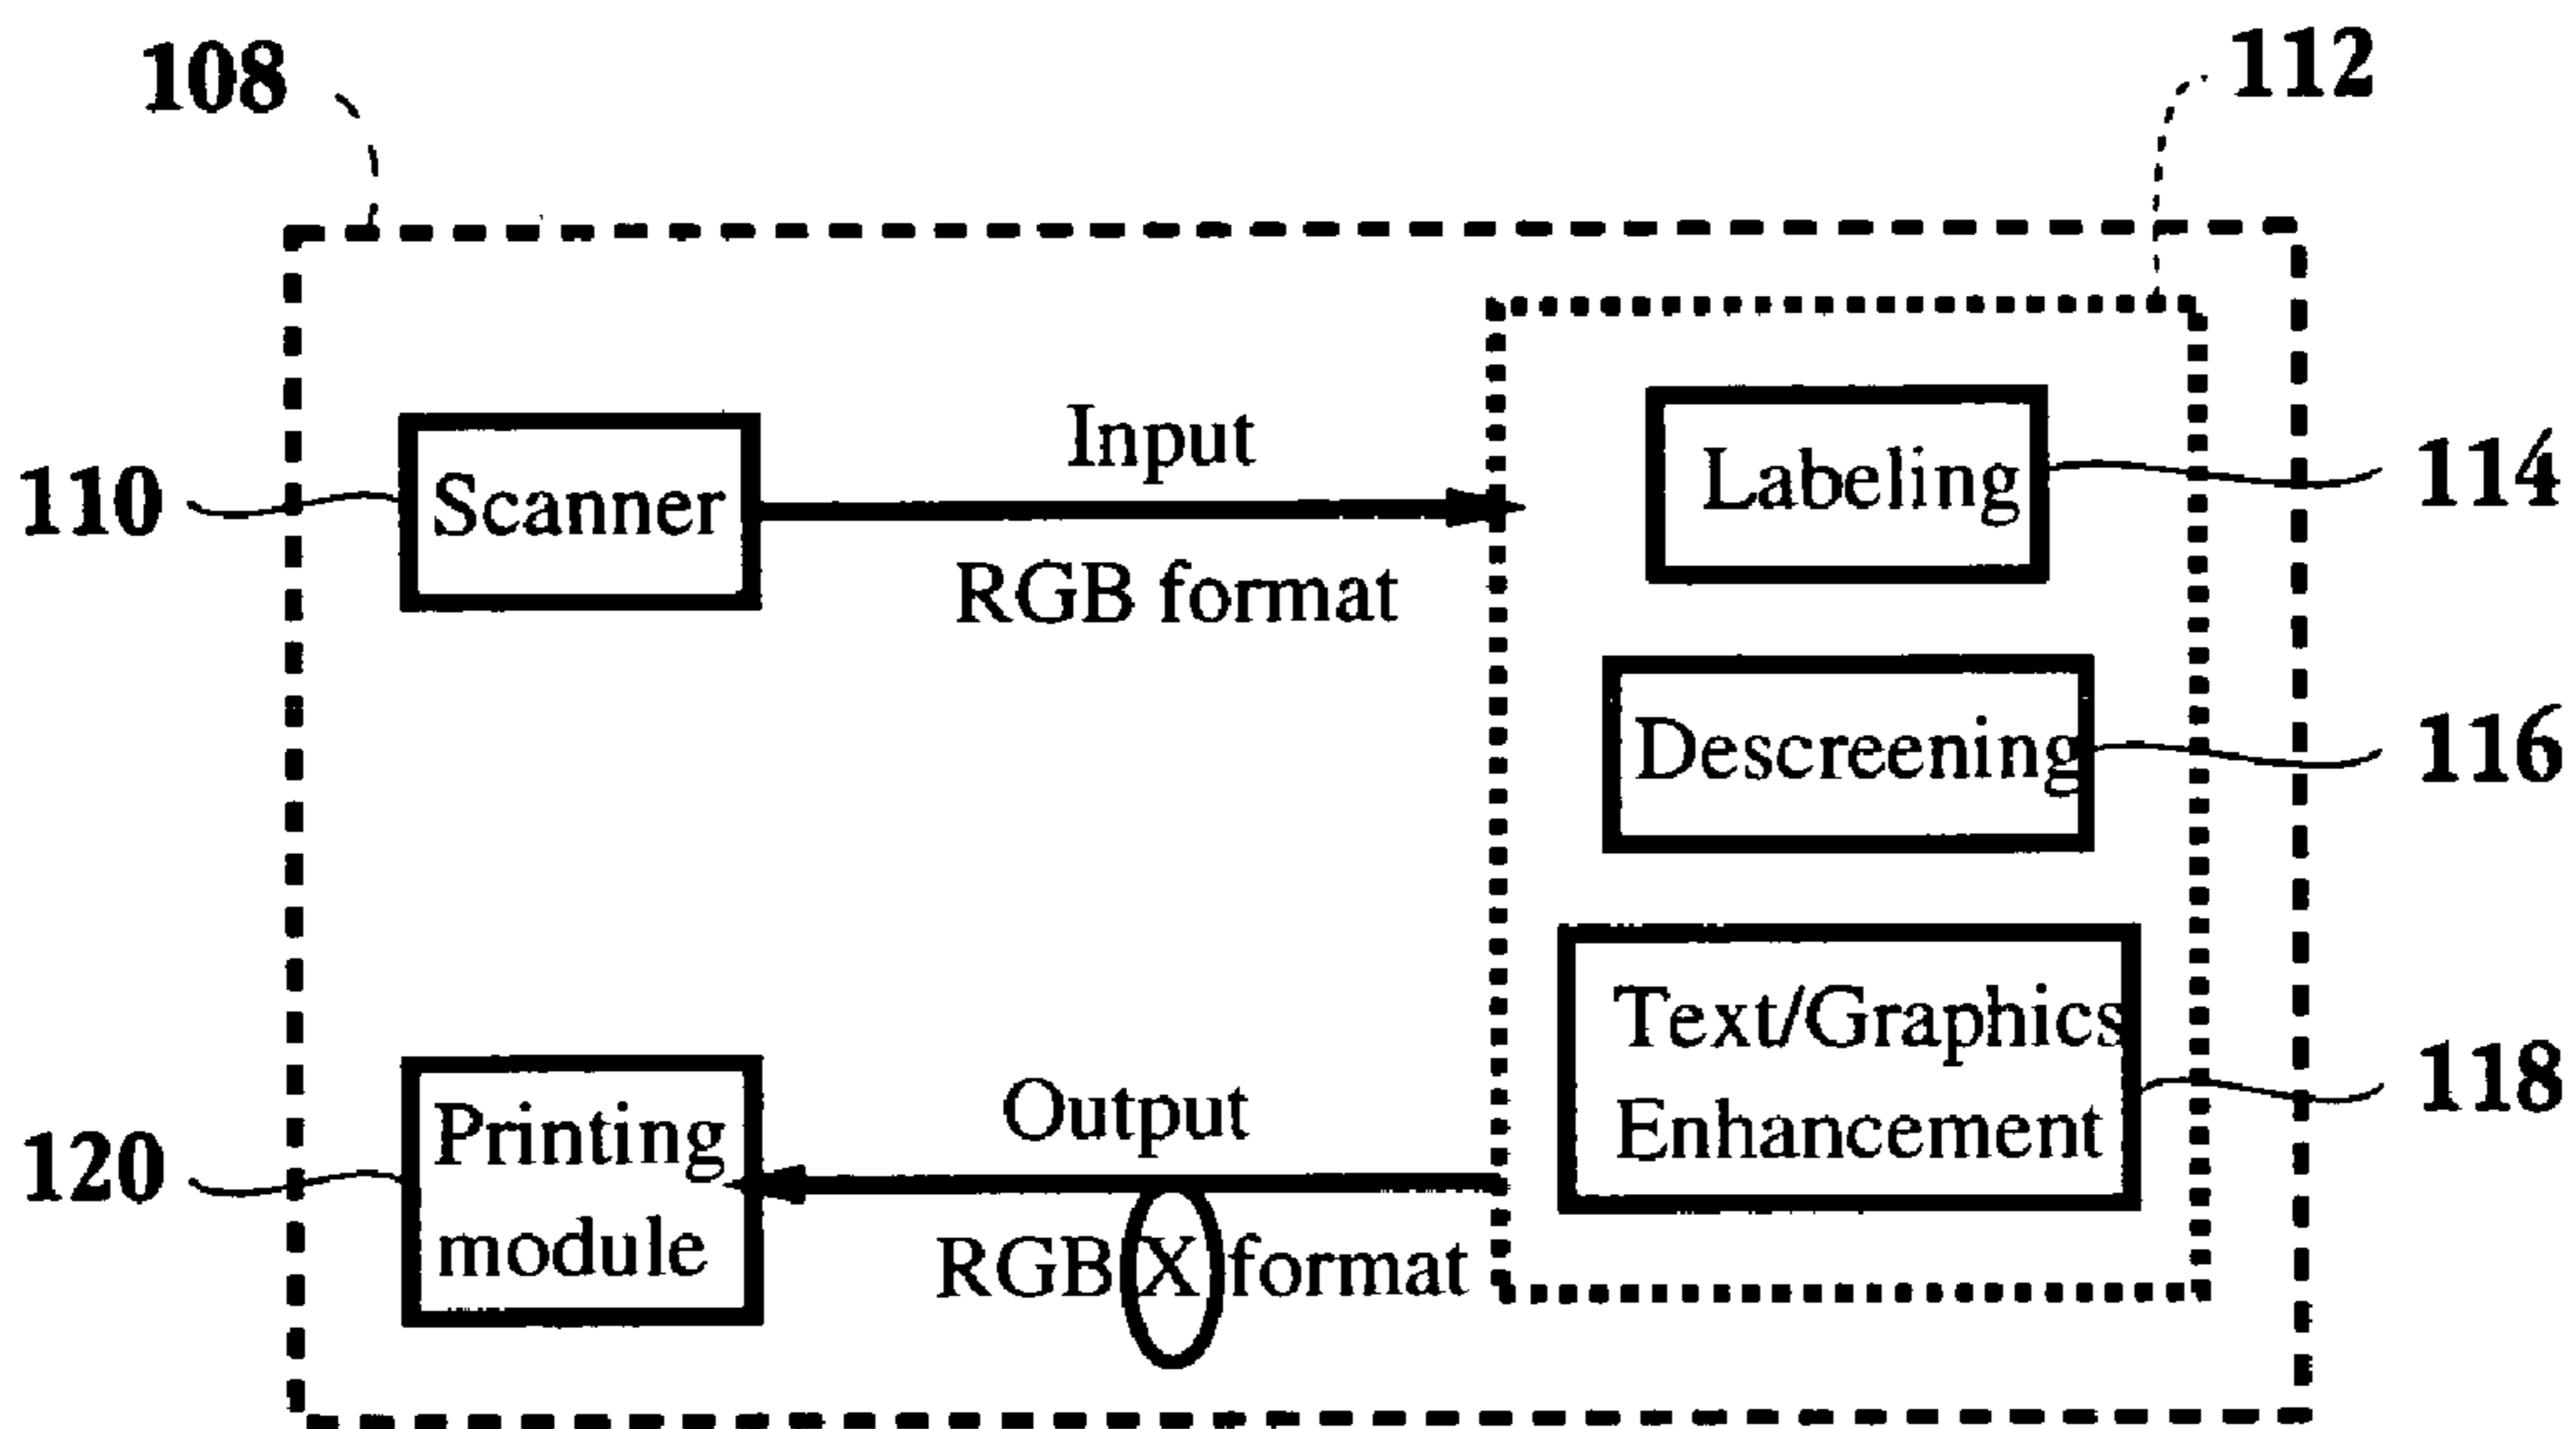


Fig. 2




Fig. 3






Fig. 4A




Fig. 4B-1




Fig. 4B-2




Fig. 4B-3




Fig. 4B-4




Fig. 4C






Fig. 5A




Fig. 5B




Fig. 6A




Fig. 6B






**Fig. 7A**




**Fig. 7B**




**Fig. 7C**

-  Thin coarse edges
-  Thick coarse edges
-  Final thin coarse edges



**Fig. 7D**



**Fig. 7E**






Fig. 8A




Fig. 8B




Fig. 9

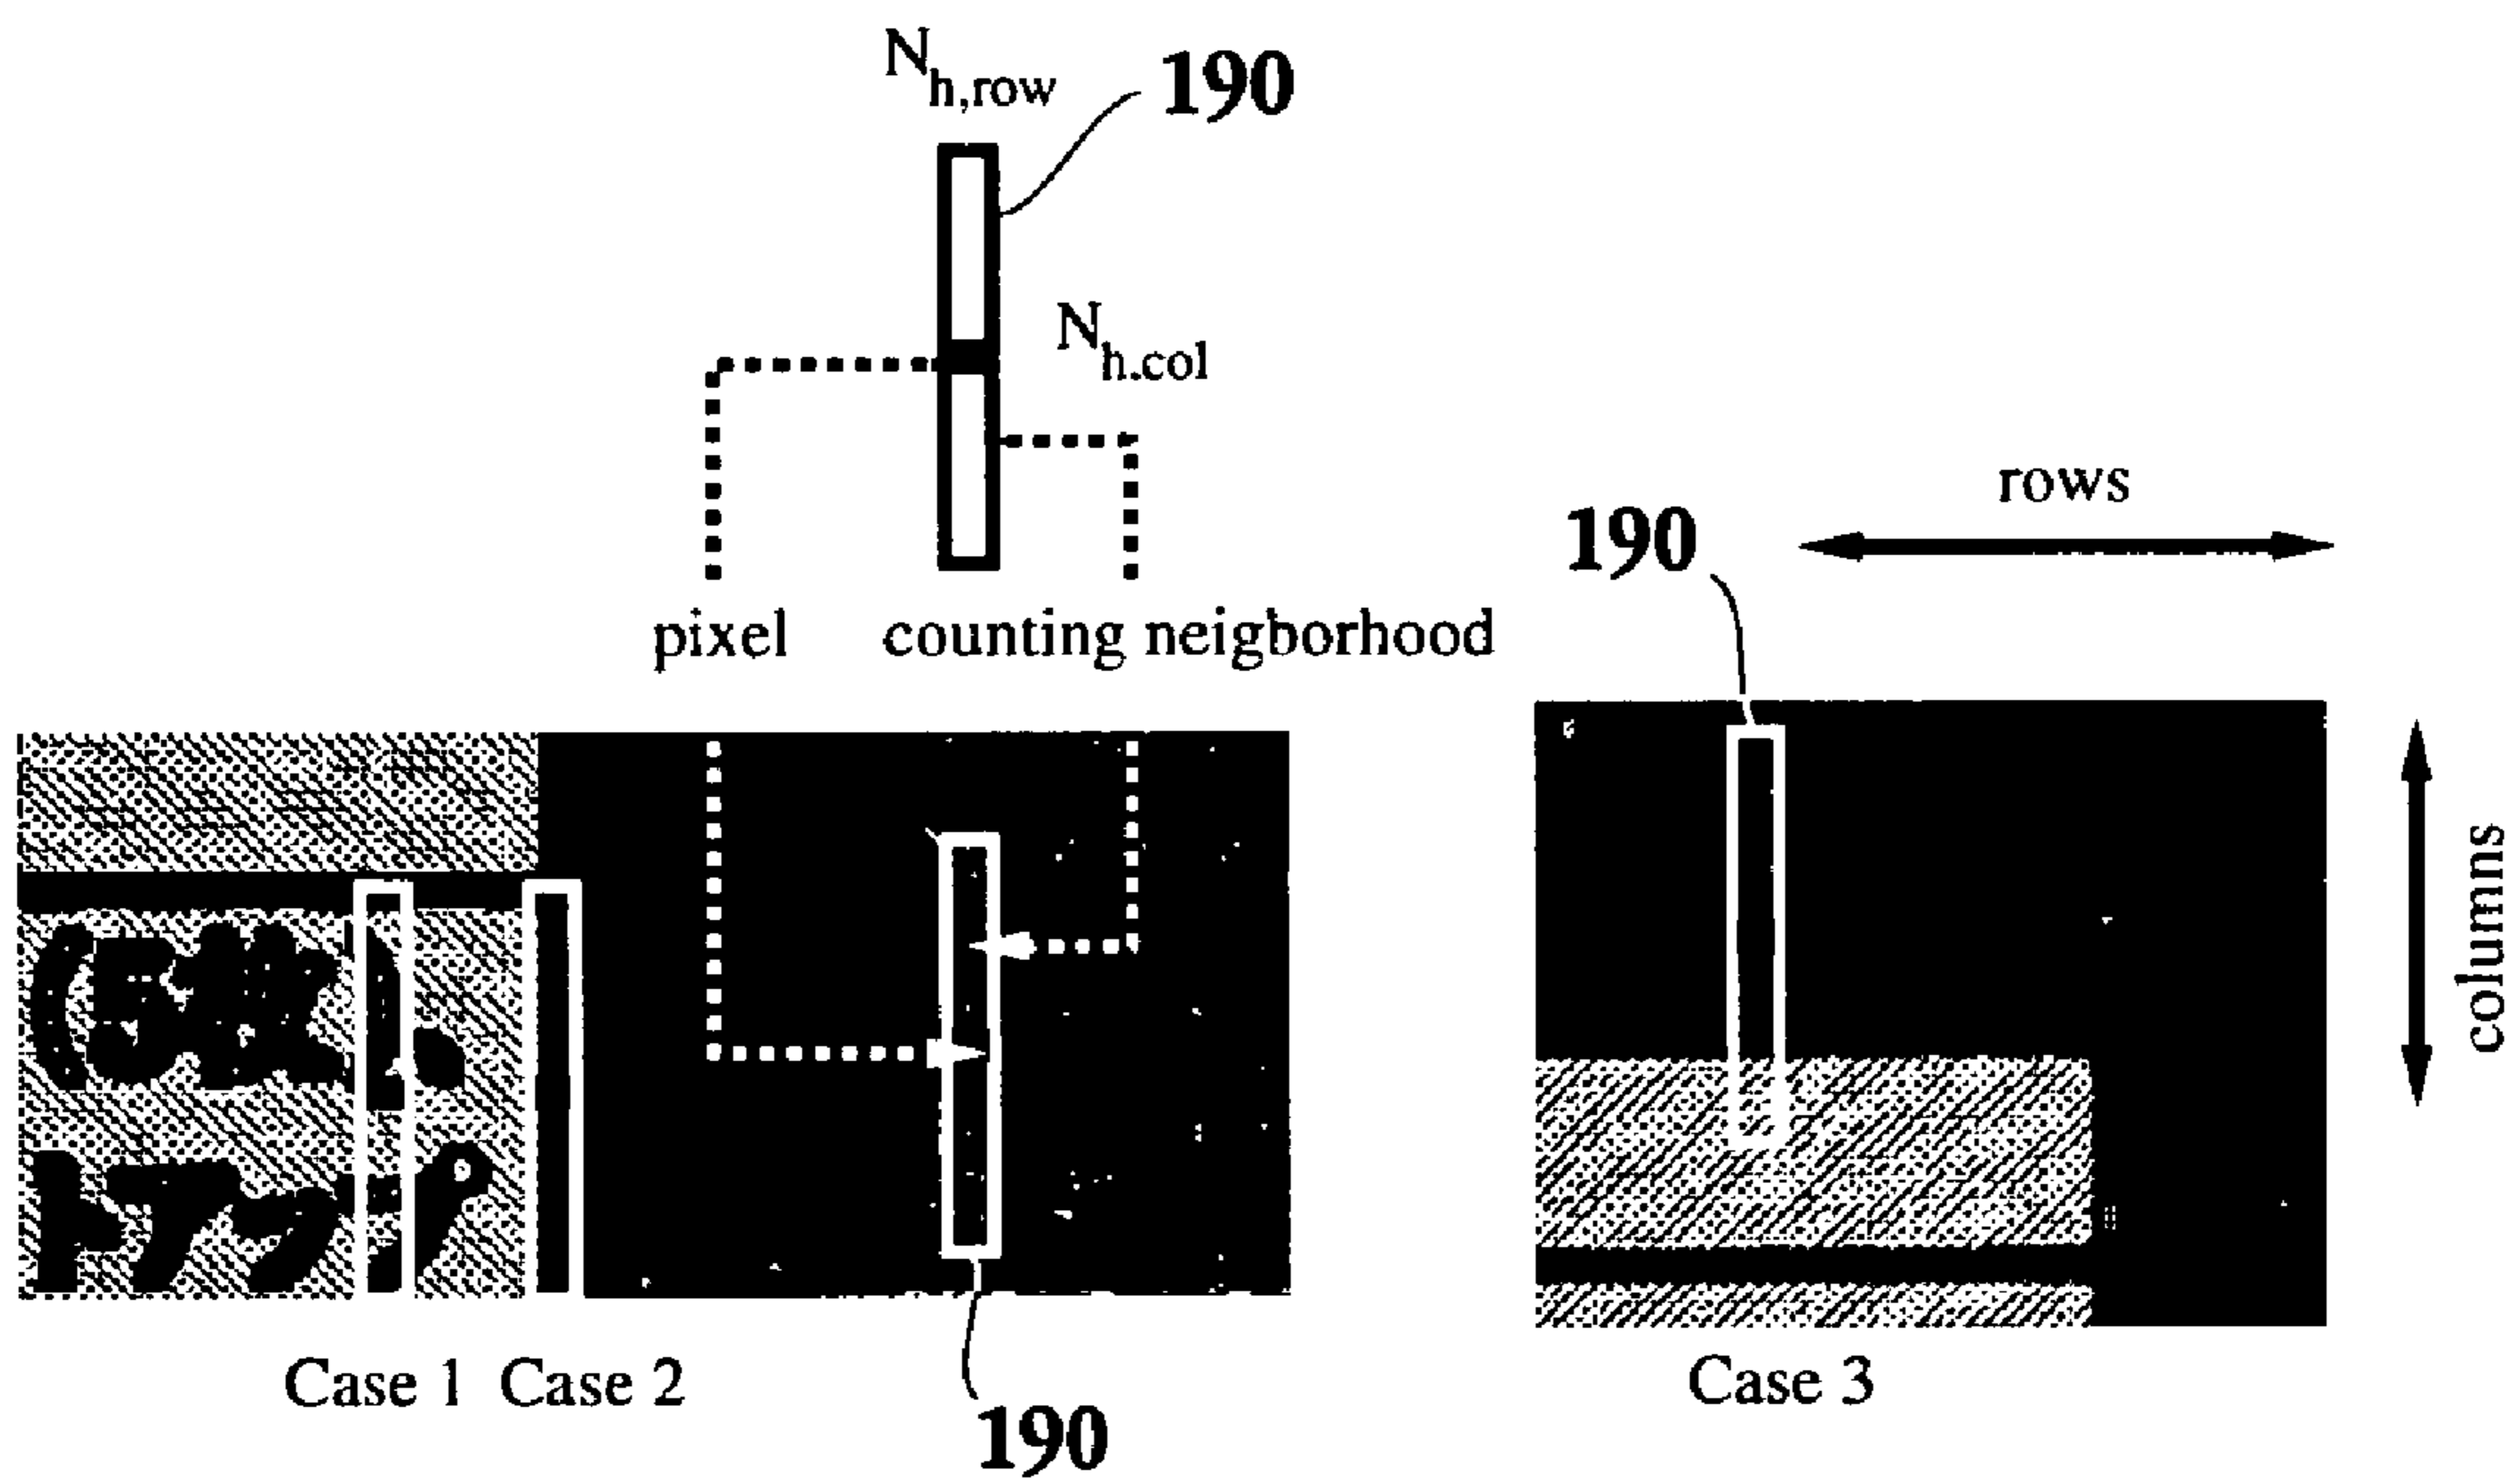




Fig. 10




**Fig. 11A**



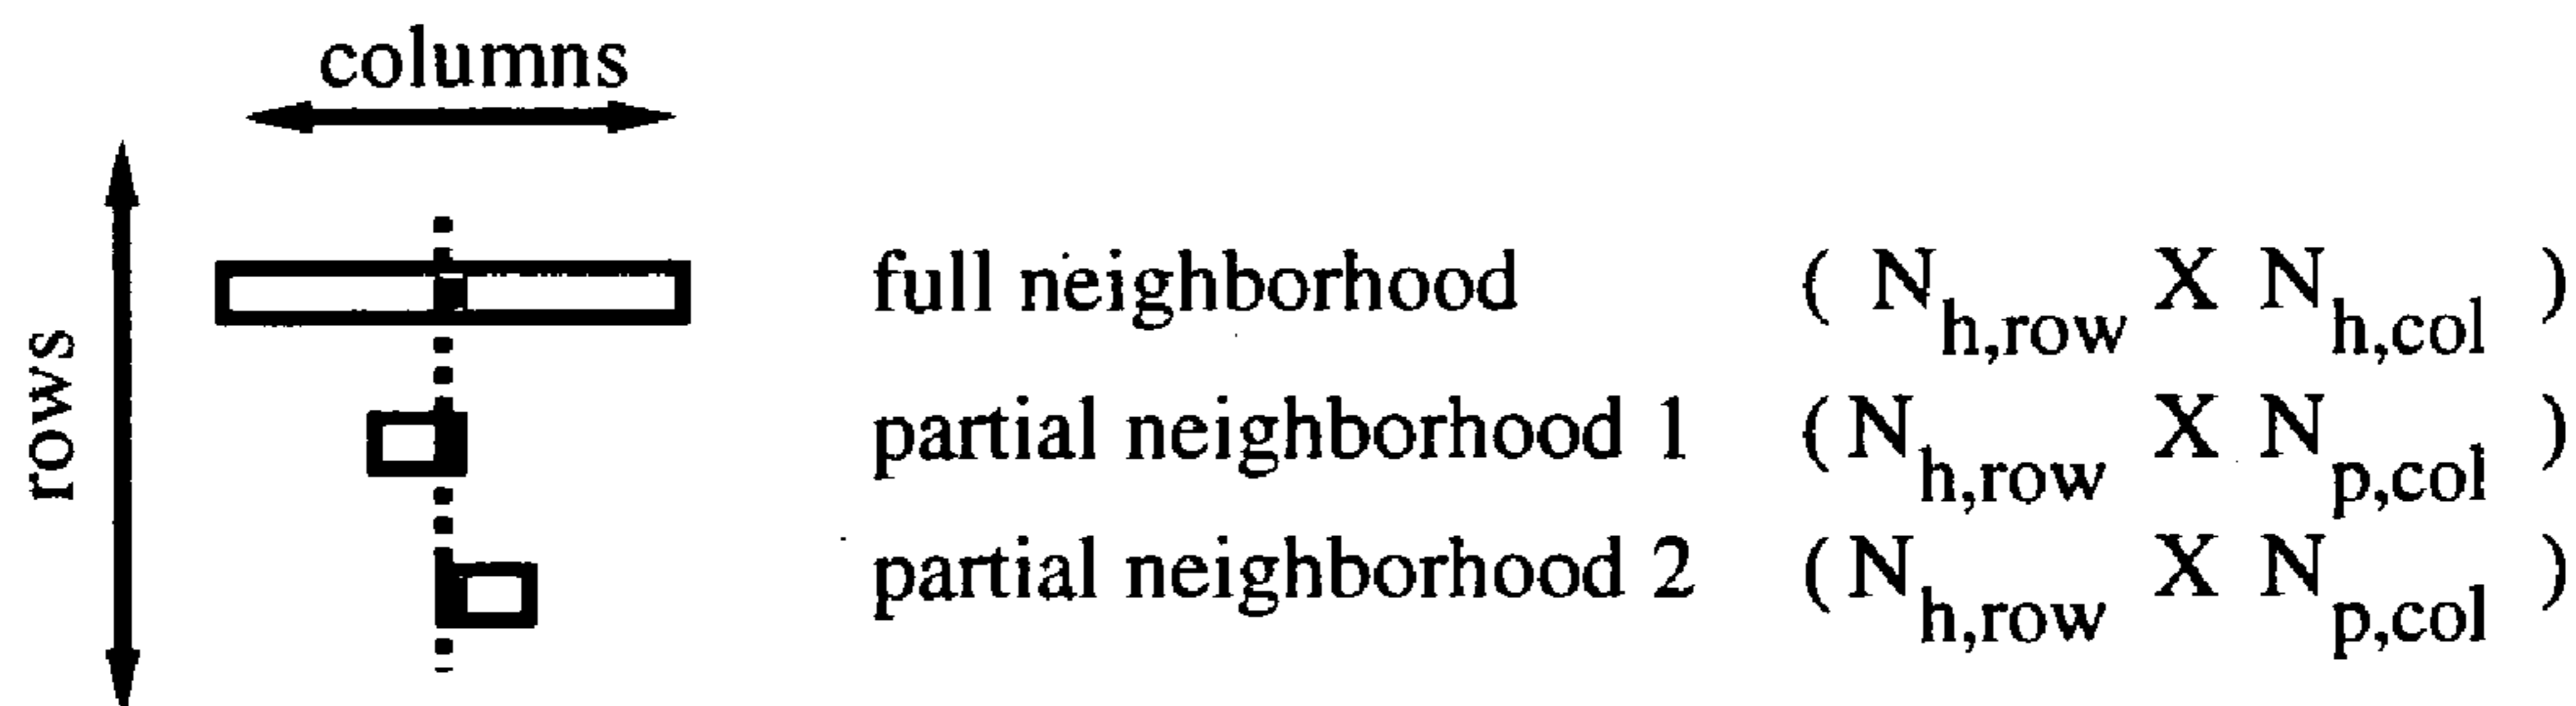
**Fig. 11B**



**Fig. 11C**



**Fig. 11D**



**Fig. 12**






Fig. 13A




Fig. 13B




Fig. 14A




Fig. 14B






Fig. 15A




Fig. 15B



仕様概要		
機種名	MJ-830C	
コントローラボード	ESC/P 7R プリンター制御ユニット搭載のプリンター	
印字方式	MJ-830C方式 (720 x 720 dpi, 360 x 360 dpi)	
紙の厚さ	0.08mm (標準)	
紙の幅	幅210mm (標準)	
印刷速度 (A4)	標準: 300dpi (10.4PPM)	
印刷速度 (A4)	高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	
印刷速度 (A4)	超高速: 300dpi (10.4PPM)	

主なオプション・消耗品		
品名	仕様	価格
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500
インク	MJ-830C	標準価格 ¥2,500

パソコン対応表 (Win98/2000は、Windows、Macintosh 専用です)		
メーカー	機種名	価格
EPSON	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500
NEC	エプソン プリンター	標準価格 ¥2,500

ソフトの対応OS一覧表		
メーカー	ソフト名	対応OS
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh
EPSON	エプソン プリンター	Windows98/2000, Macintosh

高品質を身近にするカラープリンターにディスプレイインタラクティブのCD-ROMソフト「ミックシーのはがきスタジオ」が入って新登場。

Windows98/2000/Macintosh 専用機

**MJ-830C Special2** 標準価格 ¥36,000 (税別) (本体のみ)

**MJ-510C Special2** 標準価格 ¥34,000 (税別) (本体のみ)

Example halftones

**EPSON** エプソン株式会社 〒163-8622 東京都東区南葛2-4-1 新設NSビル11階 **セイコーエプソン株式会社** 〒320-8622 長野県長野市太田1-3-5

**FAX** インフォメーション 最新カタログをはじめ、最新ドライバ情報、機種比較、Q&Aなどご購入の際に必要な最新情報がいつでも取り出せます。ファクシミリ/プッシュEメール/印刷可能機材)から電話をおかけになり、音声案内にしたがって操作して下さい。ご希望の資料はメニュー(料請求)021)で資料番号をご確認ください。

- エプソンインフォメーション 各種ドライバ/最新バージョンのFDを郵送でお届け致します。お申し込み/料金などは上記FAXインフォメーションの資料でご確認ください。
- プリンター購入ガイドインフォメーション (0424) 99-7111 受付時間9:00-17:30 月-金曜日(祝日除く) 製品の購入をお考えの方の専任窓口です。製品に関する技術的な質問など、お気軽にお電話下さい。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。
- エプソンインフォメーションセンター 製品に関するご質問/ご相談に電話でお答えします。

Fig. 16A



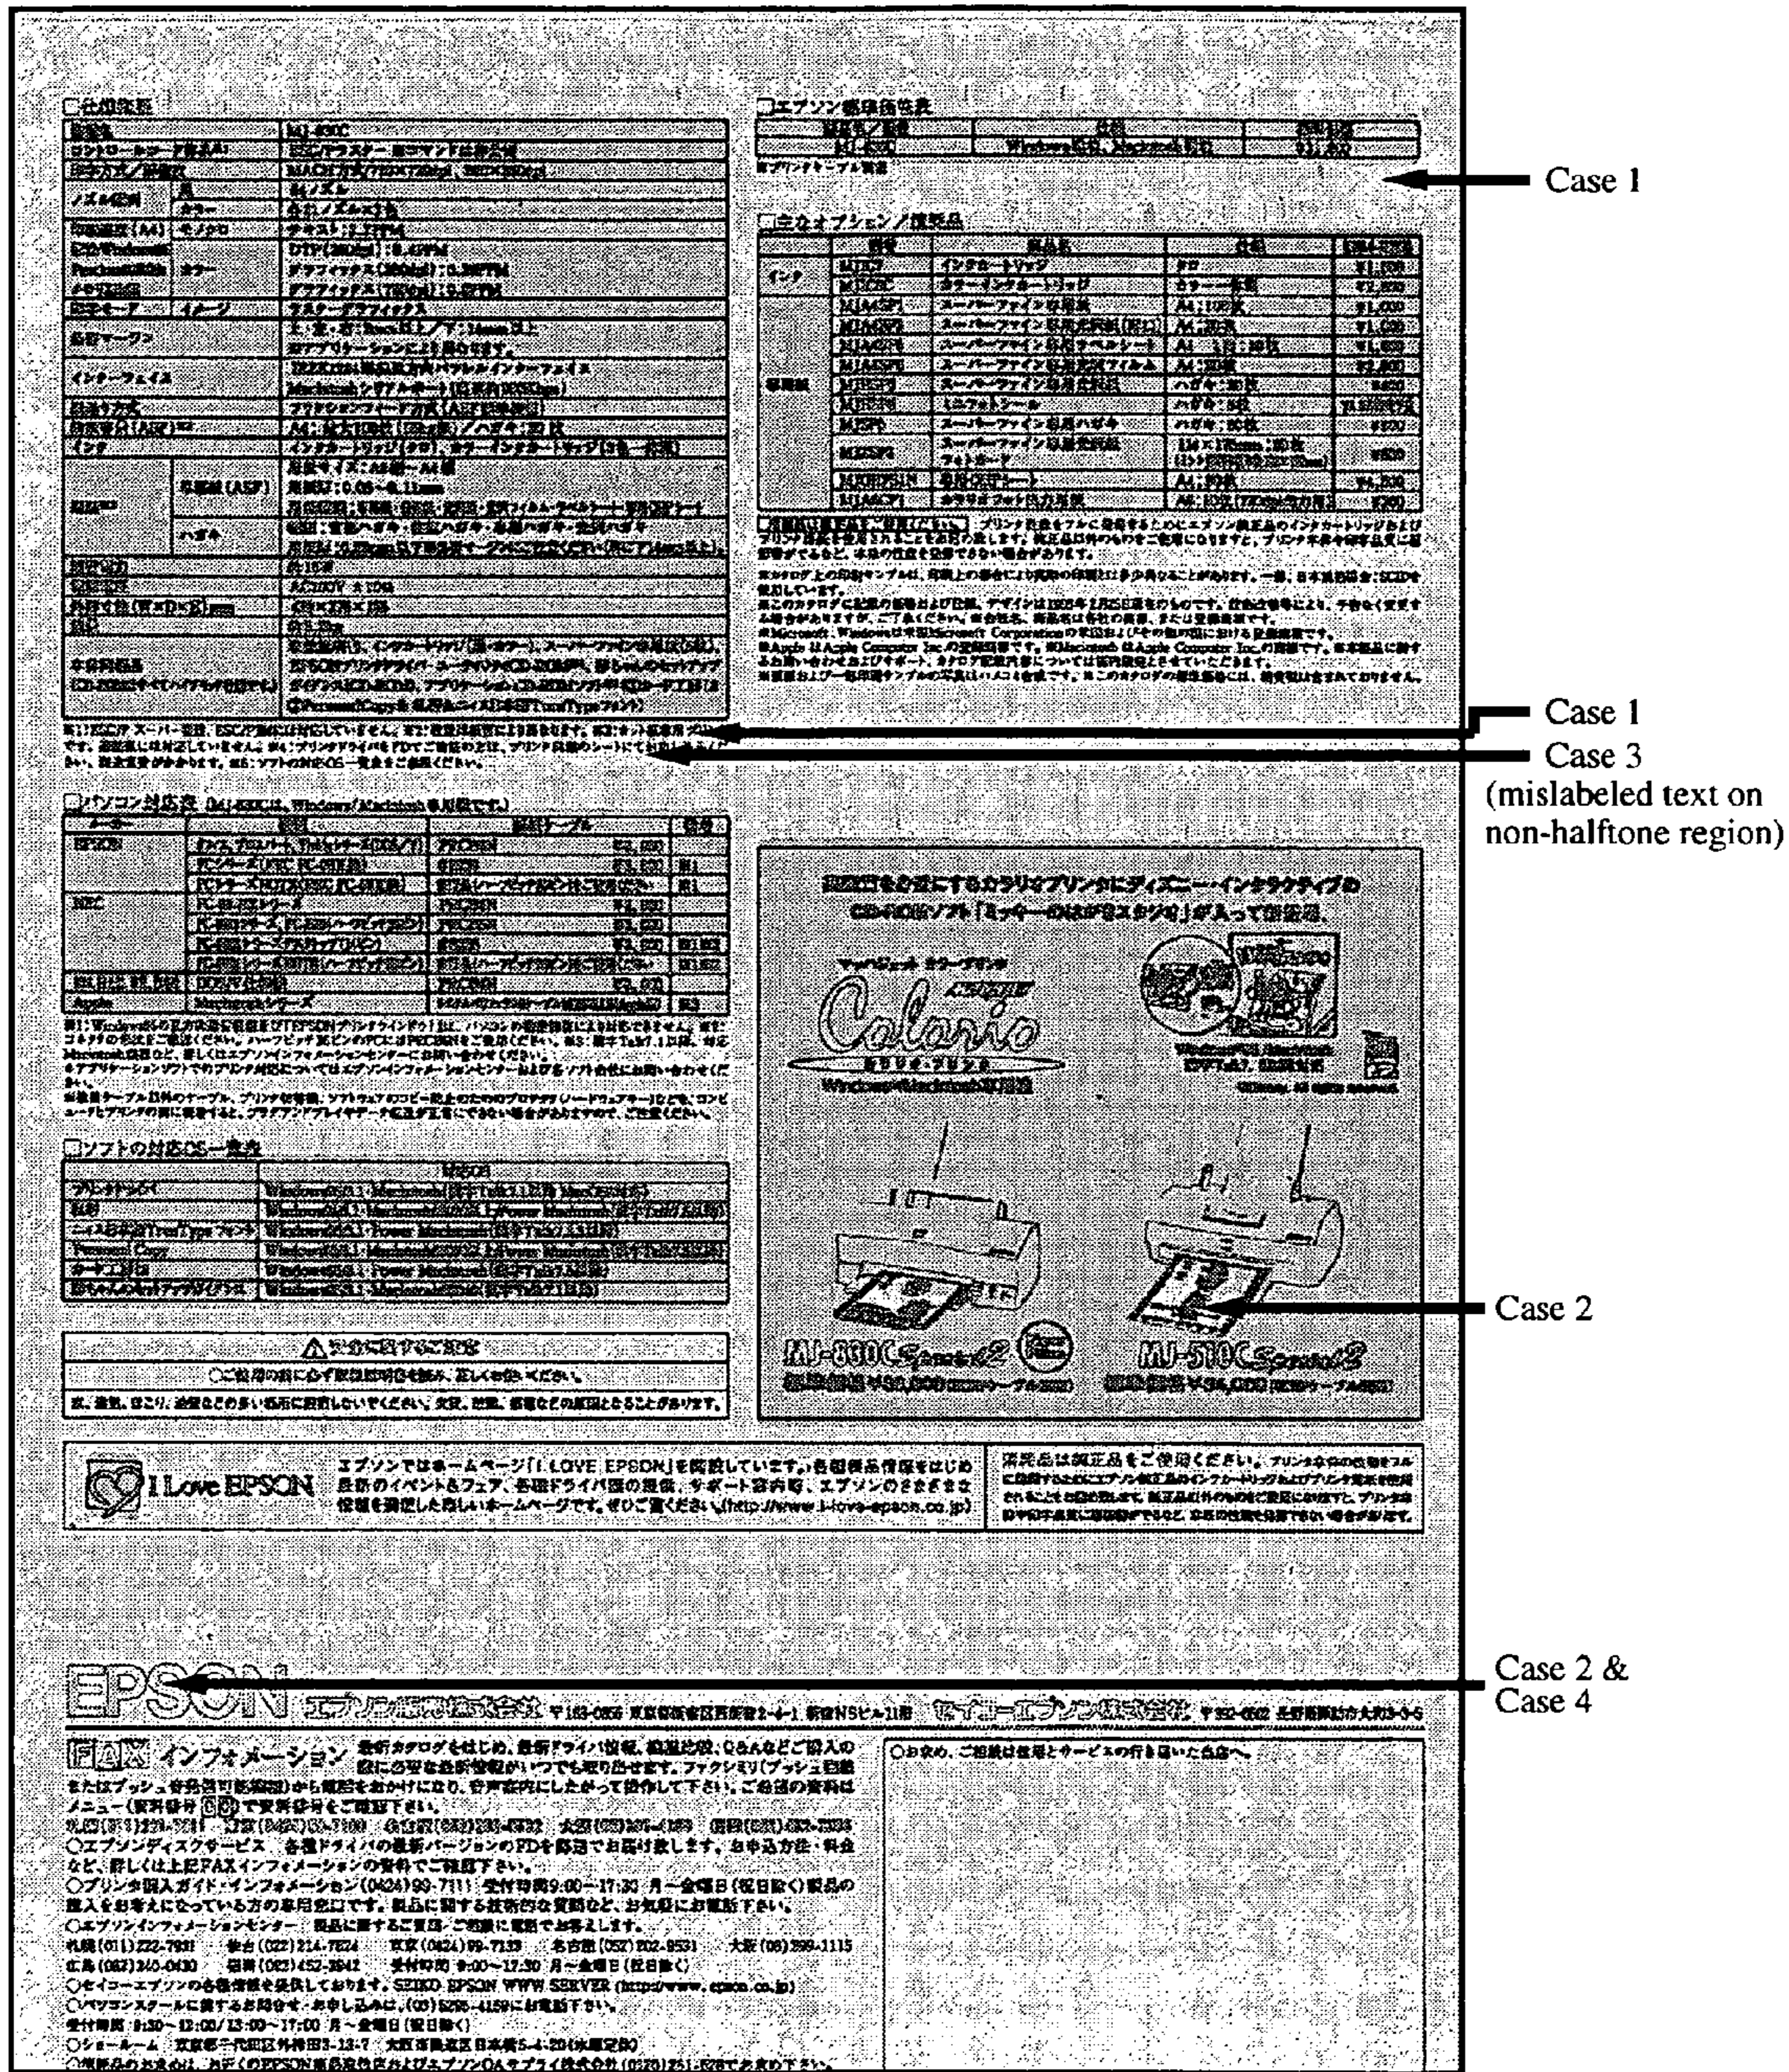



Fig. 16B

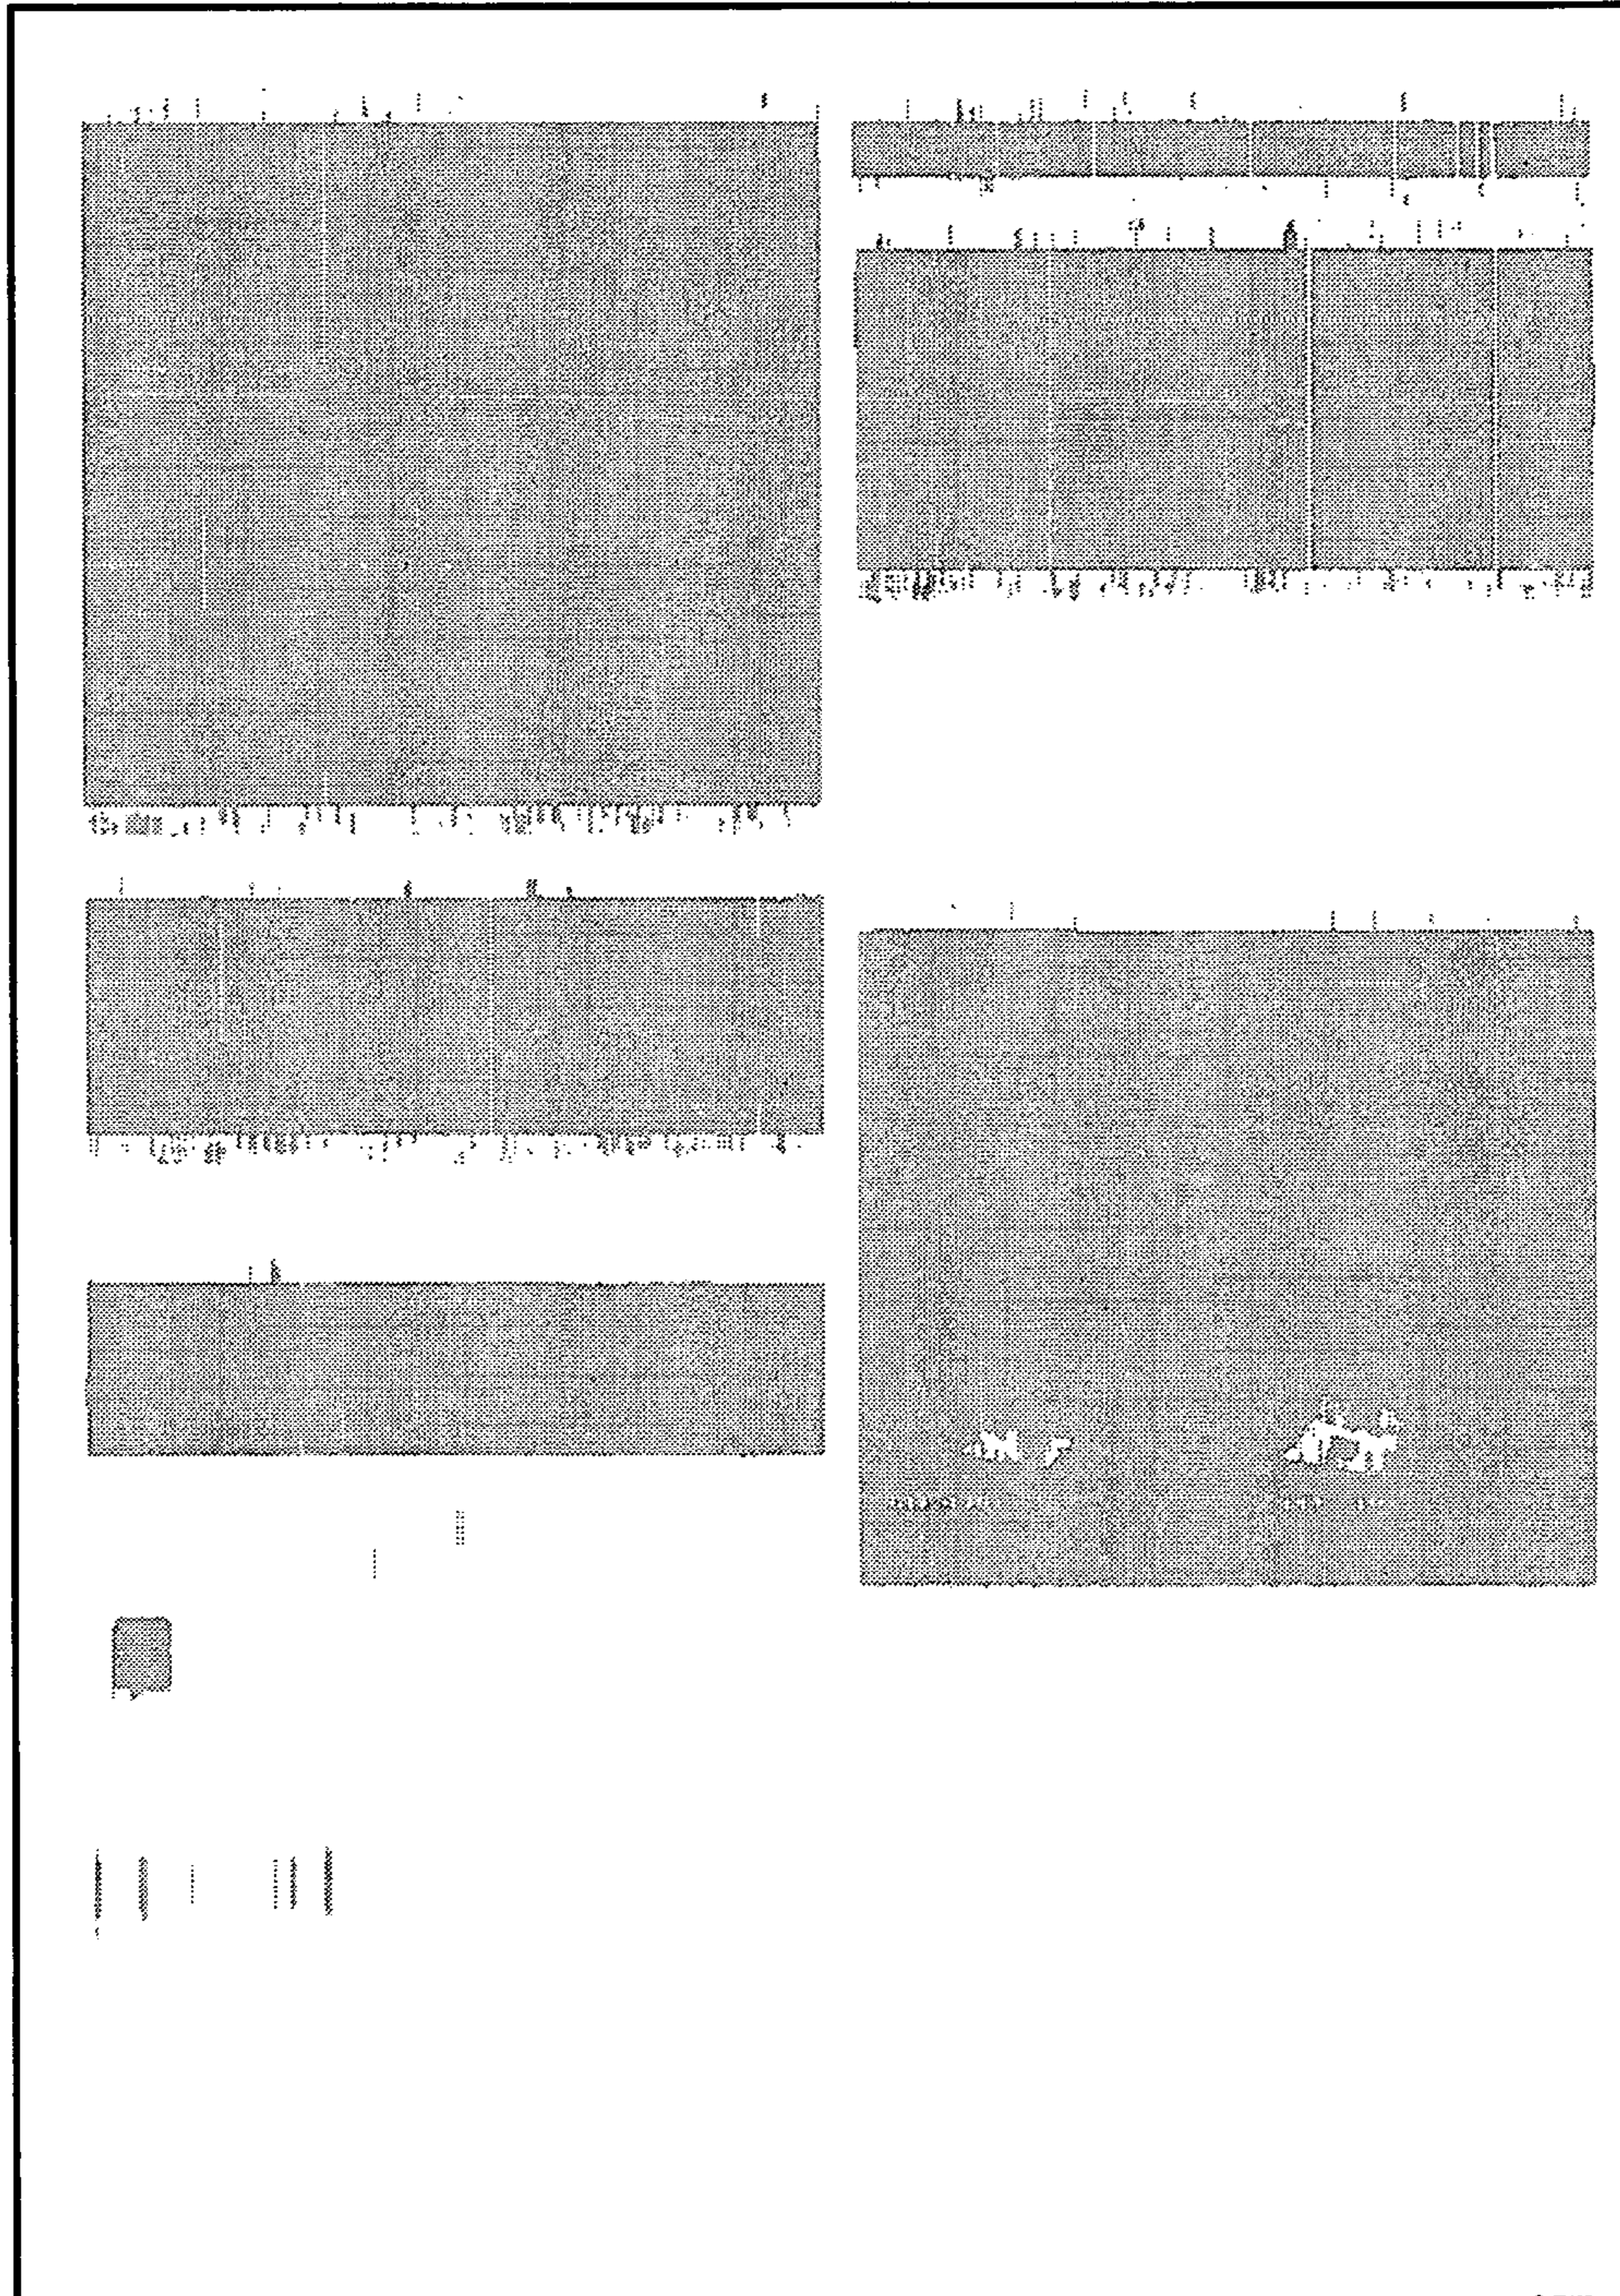




**Fig. 16C**

COLOR CODE  
■ HALFTONE  
□ NON\_HALFTONE





**Fig. 16D**

COLOR CODE  
■ HALFTONE  
□ NON\_HALFTONE



■仕様概要

Table with 2 columns: Item and Specification. Includes details on model (MJ-830C), resolution (1440x720dpi), print speed (30ppm), and power requirements.

※1: ESC/P 3.0 (スーパーファミコン) 互換性... ※2: フォントのインストール... ※3: プリンタの電源...

■パソコン対応表 (MJ-830Cは、Windows/Macintosh専用です。)

Table with 4 columns: OS, Printer Model, Resolution, and Price. Lists compatibility for Windows 95/98/Me and Macintosh systems.

※1: Windows 95の互換性... ※2: Windows 98の互換性... ※3: Macintoshの互換性...

■ソフトの対応OS一覧表

Table with 2 columns: Software Name and OS Support. Lists various software titles and their compatibility with Windows and Macintosh.

安全に関するご注意
●ご使用の前に必ず取扱説明書を読み、正しくお使いください。

■オプション標準価格表

Table with 3 columns: Item Name, Price, and Remarks. Lists optional items like ink and paper.

■主なオプション・消耗品

Table with 5 columns: Item Name, Price, and Remarks. Lists ink cartridges and other consumables.

【重要】プリンタ本体には、インクをフルに消費するためのオプション標準価格のインクカートリッジが2個付属されています。... (Note about ink cartridges and pricing)

Advertisement for the Clario printer series. Features the text '高品質を身近にするカラープリンタにディズニー・インタラクティブのCD-ROMソフト「ミッキーのはがきスタジオ」が入って新登場。' and images of the MJ-830C and MJ-510C models.

I Love EPSON logo and text: 'EPSONではホームページ「I LOVE EPSON」を開発しています。各種製品情報をはじめ最新のイベント&フェア、各種ドライバの提供、サポート案内等、EPSONのさまざまな情報を掲載した楽しいホームページです。ぜひご覧ください。' (http://www.i-love-epson.co.jp/)

EPSON エプソン販売株式会社 〒163-0655 東京都新宿区西新宿2-4-1 新宿NSビル11階 セイコーエプソン株式会社 〒222-4672 東京都大田区大田3-5-6

FAX インフォメーション section. Provides contact information for various regions: 札幌 (011) 221-7811, 東京 (0424) 88-7100, 名古屋 (052) 202-8582, 大阪 (06) 207-4358, 福岡 (092) 452-3305. Includes details on fax services and support hours.

Fig. 16E



**H**さん、今年も水仙の花をありがとう。盲導犬をつれてSNSセンターを訪れた学生さんが、「季節を感じますね」と言っていました。地域で障害者への情報提供に尽力しているあなただから、今日は少し愚痴を言わせてもらいます。さっきの学生さん、アメリカの大学を卒業して、日本で大学院に進むか、企業に就職することを望んでいるのですが、その受け入れ体制がいまひとつで……。彼は大学でインターネットを駆使して、図書館の蔵書から必要な情報入手、論文提出も一般の学生と同じかそれ以上にこなしてきました。友人と一緒に教室でイヤホン付きのノートパソコンを使って授業を受けてきたのです。情報へのアクセスは保証する。しかしそれ以上の特別扱いほしくない、という態度が当然だった世界から考えると、日本を退却のように感じるのも無理はありません。「資料は対面朗読室で読んでもらってください」「点字文書作成のボランティアを依頼します」……。好意とはいえ、その特別扱いや、情報をコントロールするのが自分自身でないことにはいらだっているのです。

残念ながら日本において、米国と同じ環境を持つ大学や企業は限られています。イントラネットが架橋に組み込まれている企業はまだ少ないし、大学図書館の電子化もこれからです。インターネットを使いたくても、スクリーンリーダーと呼ばれるGUIの読み上げソフトは、日本語処理の難しさや指導機能がないなどの問題がネックになって閲覧や操作が困難です。またLynxというテキストベースのブラウザを使って文字だけでアクセスすることも可能ですが、今度はインターネット上の絵を視覚障害者が読めるようにHTMLが書かれていないため、やはり情報が欠けてしまいます。でも、彼はここでスクリーンリーダー/2とNetscape Navigatorで、じっくりネットサーフィンしていましたよ。キーだけで操作するから私たち暗眼者よりずっと速いのです。すべての情報は、障害を持つ人にも等しく共有されるもの、という発想を、早く日本に根づかせたいですね。もちろん海外においても、まだまだ課題は残っていますが、アメリカの多くの大学が平等な情報提供をめざして参加しているEASI (Equal Access to Software and Information)のようなプロジェクトを参考に、ネットワークの力を活用していきたいと思います。日本でもインターネット上で有志が集まりはじめています。水仙が本当にいい香りです。盲導犬も、クンクンとうれしそうでしたよ。わたしと、その学生さんと、盲導犬と、一時、海浜を埋め尽くす水仙の大群落を「共有」できたのではないかと思います。ではまた！

【SNSセンターからの手紙 その2】  
関根千佳  
e17788@yamato.ibm.co.jp

\*1 障害者が快適にWWWを利用するためのホームページ (Lynxがダウンロード可能)  
<http://www.ibm.co.jp/kokoroweb/tips/>

\*2 EASI <http://www.rit.edu/~easi/>

■ 編集プロフィール  
せまはるか  
日本IBM SNS (スペシャル・ニーズ・システム) センター主任。  
パソコン雑誌「People」(雑誌「ワークラップ」)プロデューサーでもあり(ハンドル名:はるか)。長崎県佐賀市生まれ。

■ 日本IBM SNSセンター  
障害者向けのパソコン利用を支援するセンター。  
東京・調布にある。  
このコラムはSNSセンターでの実話に基づいている。

1997.2 Intranet 205

ゆにはばーさる  
あくせすの研究

スクリーンリーダーが  
広げる世界





Fig. 17A



**H**さん、今年も水仙の花をありがとう。盲導犬をつれてセンターを訪れた学生さんが、「季節を感じますね」と言っていました。

地域で障害者への情報提供に尽力しているあなただから、今日は少し感動を味わわせてもらいます。さっきの学生さん、アメリカの大学を卒業して、日本で大学院に進むか、企業に就職することを望んでいるのですが、その受け入れ体制がいまひとつで……。彼は大学でインターネットを駆使して、図書館の検索から必要な情報入手、論文提出も一般の学生と同じかそれ以上にこなしてきました。友人と一緒に教室でイヤホン付きのノートパソコンを使って授業を受けてきたのです。情報へのアクセスは保証する、しかしそれ以上の特別扱いはいらない、という態度が当然だった世界から考えると、日本も遠征のように感じるのも無理はありません。「資料は対面紙読書で読んでもらってください」「点字文書作成のボランティアを依頼します」……。好意とはいえ、その特別扱いや、情報をコントロールするのが自分自身でないことにはらだっているのです。

残念ながら日本において、米国と同じ環境を持つ大学や企業は限られています。イントラネットが業務に組み込まれている企業はまだ少ないし、大学図書館の電子化もこれからです。インターネットを使いたくても、スクリーンリーダーと呼ばれるGUIの読み上げソフトは、日本語処理の難しさや相違機がないなどの課題がネックになって開発や普及が困難です。またLynxというテキストベースのブラウザを使って文字だけでアクセスすることも可能ですが、今度はインターネット上の絵を視覚障害者が読めるようにHTMLが整っていないため、やはり情報が欠けてしまいます。でも、彼はこのスクリーンリーダー/2とNetscape Navigatorで、しっかりネットサーフィンしていましたよ。キーだけで操作するから私たち障害者よりずっと速いのです！

すべての情報は、障害を持つ人にも等しく共有されるもの、という発想を、早く日本に根づかせたいですね。もちろん海外においても、まだまだ問題は残っていますが、アメリカの多くの大学が平等な情報提供をめざして参加しているEASI (Equal Access to Software and Information)<sup>1,2</sup>のようなプロジェクトを参考に、ネットワークの力を活用していきたいと思います。日本でもインターネット上で有志が集まりはじめています。水仙が本当にいい香りです。盲導犬も、ケンケンとうれしそうでしたよ。わたしと、その学生さんと、盲導犬と、一瞬、海岸を埋め尽くす水仙の大群落を「共有」できたのではないかと感じます。ではまた！

SN3センターからの手紙 その2  
 岡原千佳  
 e17788@yamato.ibm.co.jp

1 障害者が快適にWWWを利用するためのホームページ (Lynxがダウンロード可能)  
<http://www.ibm.co.jp/kokoroweb/t/pa/>  
 \*2 EASI <http://www.rit.edu/~easi/>

スクリーンリーダーが  
 広げる世界

読者プロフィール  
 せきねわか  
 日本IBM SMS (スペシャル・ニーズ・システム) センター主任  
 パソコン設置People's「福祉工作クラブ」プロデューサーでもある (ハンドル名: さえら)。  
 長崎県佐賀県在住。


1. 日本IBM SMS センター  
 障害を持つ人のパソコン利用を支援するセンター。  
 以下：福岡にある。  
 このコラムは同センターでの  
 取材を基にしている。

Case 4  
 (missed small  
 halftone region)

Fig. 17B

- COLOR CODE
- FINE EDGE
  - COARSE EDGE
  - HALFTONE
  - CONSTANT TONE
  - OTHER






**Fig. 17C**

COLOR CODE  
■ HALFTONE  
□ NON\_HALFTONE





**Fig. 17D**

COLOR CODE  
■ HALFTONE  
□ NON\_HALFTONE



**H**さん、今年も水仙の花をありがとう。盲導犬をつれてセンターを訪れた学生さんが、「季節を感じますね」と言っていました。地域で障害者への情報提供に尽力していらっしゃるあなただから、今日は少し愚痴を言わせてもらいます。さっきの学生さん、アメリカの大学を卒業して、日本で大学院に進むか、企業に就職することを望んでいるのですが、その受け入れ体制がいまひとつで……。彼は大学でインターネットを駆使して、図書館の検索から必要な情報入手、論文提出も一般の学生と同じかそれ以上にこなしてきました。友人と一緒に教室でイヤホン付きのノートパソコンを使って授業を受けてきたのです。情報へのアクセスは保証する、しかしそれ以上の特別扱いはいらない、という態度が当然だった世界から考えると、日本を遺跡のように感じるのも無理はありません。「資料は対面読書室で読んでもらってください」「点字文書作成のボランティアを依頼します」……。好意とはいえ、その特別扱いや、情報をコントロールするのが自分自身でないことにいらだっているのです。

残念ながら日本において、米国と同じ環境を持てる大学や企業は限られています。イントラネットが業務に組み込まれている企業はまだ少ないし、大学図書館の電子化もこれからです。イン

ターネットを使いたくても、スクリーンリーダーと呼ばれるGUIの読み上げソフトは、日本語処理の難しさや指導機関がないなどの問題がネックになって開発や普及が困難です。またLynxというテキストベースのブラウザを使って文字だけでアクセスすることも可能ですが、今度はインターネット上の絵を視覚障害者が読めるようにHTMLが書かれていないため<sup>\*1</sup>、やはり情報が欠けてしまいます。でも、彼はここでスク

リーンリーダー/2とNetscape Navigatorで、しっかりネットサーフィンしていましたよ。キーだけで操作するから私たち暗眼者よりずっと速いのです。すべての情報は、障害を持つ人にも等しく共有されるもの、という発想を、早く日本に根づかせたいですね。もちろん海外においても、まだまだ課題は残っていますが、アメリカの多くの大学が平等な情報提供をめざして参加しているEASI (Equal Access to Software and Information)<sup>\*2</sup>のようなプロジェクトを参考に、ネットワークの力を活用していきたいと思います。日本でもインターネット上で有志が集まりはじめています。水仙が本当にいい香りです。盲導犬も、ククンとうれしそうでしたよ。わたしと、その学生さんと、盲導犬と、一瞬、海岸を埋め尽くす水仙の大群落を「共有」できたのではないかと思います。ではまた!

【SNSセンターからの手紙 その2】

関根千佳

ei7788@yamato.ibm.co.jp

\*1 障害者が快適にWWWを利用するためのホームページ (Lynxがダウンロード可能)

<http://www.ibm.co.jp/kokoroweb/tips/>

\*2 EASI <http://www.rit.edu/~easi/>



ゆにばーさる  
あくせすの研究

# スクリーンリーダーが 広げる世界

著者プロフィール  
せきねちか  
日本IBM SNS (スペシャル・ニーズ・システム) センター主任。  
パソコン通信People[福祉工作クラブ]プロデューサーでもあり(ハンドル名:せきねちか)。  
長崎県佐賀市生まれ。

IBM SNSセンター  
障害者特製のパソコン利用を  
支援するセンター。  
東京(有明)にある。  
このコラムはSNSセンターでの  
取材を基にしている。

Fig. 17E



"we did not do anything that is related to financial structuring or disclosure or any of the issues that got them into trouble. We stand by all the work we did. Beyond that, we can only empathize with the trouble they are going through. It's a sad thing to see."

Still, outsiders marvel that the secretive partnership has not been drawn into the debacle, given its extensive involvement at Enron. "I'm surprised that they haven't been subpoenaed as a witness, at least," says Wayne E. Cooper, CEO of Kennedy Information, a research and publishing firm that keeps tabs on consultants. "There was so much smoke coming out of the Andersen smoking gun that all the firefighters went after that one. McKinsey was lucky. They dodged a bullet."

The bad news, however, is that Enron, which was

Did McKinsey's partners get caught up in the euphoria of the late '90s and suffer lapses of judgment? And if so, what does that say about the quality of its expensive advice? Did it stray from its core values? What accountability does it—or any consulting firm—have for the ideas and concepts it launches into a company?

After all, McKinsey was a key architect of the strategic thinking that made Enron a Wall Street darling. In books, articles, and essays, its partners regularly stamped their imprimatur on many of Enron's strategies and practices, helping to position the energy giant as a corporate innovator worthy of emulation. The firm may not be the subject of any investigations, but its close involvement with Enron raises the question of whether McKinsey, like some other professional firms, ignored warning flags in order to keep an important account.

The breakdowns of such visible clients could not have come at a more trying time. Instead of celebrating the end of his third and final three-year term

**MID-1970s** New strategy consulting boutiques Boston Consulting Group and Bain offer first intense competition in high-level strategy consulting.

**1978** McKinsey partner Lou Gerstner departs firm for American Express and later becomes CEO of RJR Nabisco and IBM. Other alumni head some of the world's biggest companies.

**1984** Launches *The McKinsey Quarterly*, which remains the foremost journal of management ideas produced by a consulting firm.

as managing director, Gupta, 53, finds his firm rolled by a rare and potentially disruptive downturn in its business. Like most other consulting firms, McKinsey rode the e-business wave to record revenues—and record partner payouts—in 2000. When the boom turned to bust, the firm was stuck with far too many consultants and not nearly enough assignments. The utilization rate, or billable time, of its consultants has fallen to its lowest level in more than 32 years: just 62%, vs. the heady 64% level during the dot-com boom.

That's not to say that McKinsey has lost its standing. The firm remains the high priest of high-level consulting, with the

**McKINSEY OVER THE YEARS**

**1933** Marvin Bower, a lawyer and Harvard MBA, joins the New York office and essentially invents the discipline of modern management consulting.

**1935** Marshall Field & Co becomes the firm's first big client and eventually hires James McKinsey to be CEO.

**1926** University of Chicago accounting professor James "Mac" McKinsey founds the firm. First client: meatpacker Armour & Co.

paying McKinsey as much as \$10 million in annual fees, is just one of an unusual number of embarrassing client failures for the elite consulting firm. Besides Enron, there's Swissair, Kmart, and Global Crossing—all McKinsey clients that have filed for bankruptcy in relatively short order. And those are just the biggest. McKinsey also finds itself improbably lining up with other creditors to collect unpaid fees from recently bankrupt companies that soared during the late '90s only to crash later. Battery maker Exide Technologies and NorthPoint Communications Group Inc., an upstart telecom provider, are two such examples.

All of which raises uncomfortable questions about the world's most prestigious—and enigmatic—consulting firm.

**1959** Opens first non-U.S. office in London and wins landmark assignment with Bank of England.

**1960s** Successful expansion overseas throughout Europe and Australia.

**1976** Launches *The McKinsey Quarterly*, which remains the foremost journal of management ideas produced by a consulting firm.

**1984** Launches *The McKinsey Quarterly*, which remains the foremost journal of management ideas produced by a consulting firm.

**1990s** Enron becomes a major client, leading to a record year for the firm in 2000.

**2000** Enron's collapse leads to a sharp decline in McKinsey's utilization rate.

**2001** McKinsey's utilization rate falls to its lowest level in 32 years.

**2002** McKinsey's utilization rate remains low.

**2003** McKinsey's utilization rate remains low.

**2004** McKinsey's utilization rate remains low.

**2005** McKinsey's utilization rate remains low.

**2006** McKinsey's utilization rate remains low.

Example halftones






Fig. 18A



"we did not do anything that is related to financial structuring or disclosure or any of the issues that got them into trouble. We stand by all the work we did. Beyond that, we can only empathize with the troubles they are going through. It's a sad thing to see."

Bill, outsiders marvel that the secretive partnership has not been drawn into the debate, given its extensive involvement at Enron. "I'm surprised that they haven't been subpoenaed as a witness, at least," says Wayne E. Cooper, CEO of Kennedy Information, a research and publishing firm that keeps tabs on consultants. "There was so much smoke coming out of the Andersen smoking gun that all the firefighters went after that one. McKinsey was lucky. They dodged a bullet."

The bad news, however, is that Enron, which was

Did McKinsey's partners get caught up in the euphoria of the late '90s and suffer lapses of judgment? And if so, what does that say about the quality of its expensive advice? Did it stray from its core values? What accountability does it—or any consulting firm—have for the ideas and concepts it launches into a company?

After all, McKinsey was a key architect of the strategic thinking that made Enron a Wall Street darling. In books, articles, and essays, its partners regularly stamped their imprimatur on many of Enron's strategies and practices, helping to position the energy giant as a corporate innovator worthy of emulation. The firm may not be the subject of any investigations, but its close involvement with Enron raises the question of whether McKinsey, like some other professional firms, ignored warning signs in order to keep an important account.

The breakdowns of such visible clients could not have come at a more trying time. Instead of celebrating the end of his third and final three-year term

**1989-1997** New strategy consulting boutiques Boston Consulting Group and Bain offer first intense competition in high-level strategy consulting.

**1978** McKinsey partner Lou Gerstner departs firm for American Express and later becomes CEO of RJR Nabisco and IBM. Other alumni head some of the world's biggest companies.

**1990** Launches *The McKinsey Quarterly*, which remains the foremost journal of management ideas produced by a consulting firm.

**1999** Opens first non-U.S. office in London and wins landmark assignment with Bank of England.

**1990s** Successful expansion overseas throughout Europe and Australia.

As managing director, Gupta, 58, finds his firm rolled by a rare and potentially disruptive downturn in its business. Like most other consulting firms, McKinsey rode the e-business wave to record revenues—and record partner payouts—in 2000. When the boom turned to bust, the firm was stuck with far too many consultants and not nearly enough assignments. The utilization rate, or billable time, of its consultants has fallen to its lowest level in more than 32 years: just 67%, vs. the heady 94% level during the dot-com boom.

That's not to say that McKinsey has lost its standing. The firm remains the high priest of high-level consulting, with the

**1953** Marvin Desser, a lawyer and Harvard MBA, joins the New York office and essentially invents the discipline of modern management consulting.

**1955** Marshall Field & Co. becomes the firm's first big client and eventually hires James McKinsey to be CEO.

**1996** University of Chicago accounting professor James "Mac" McKinsey founds the firm. First client: meatpacker Armour & Co.

paying McKinsey as much as \$10 million in annual fees, is just one of an unusual number of embarrassing client failures for the elite consulting firm. Besides Enron, there's Swissair, Kmart, and Global Crossing—all McKinsey clients that have filed for bankruptcy in relatively short order. And those are just the biggest. McKinsey also finds itself improbably lining up with other creditors to collect unpaid fees from recently bankrupt companies that occurred during the late '90s only to crash later. Battery maker Exide Technologies and NorthPoint Communications Group Inc., an upstart telecom provider, are two such examples.

All of which raises uncomfortable questions about the world's most prestigious—and enigmatic—consulting firm.


**McKINSEY OVER THE YEARS**

Case 5 (transition from halftone to non-halftone)

Fig. 18B

- COLOR CODE
- FINE EDGE
  - COARSE EDGE
  - HALFTONE
  - CONSTANT TONE
  - OTHER






**Fig. 18C**

COLOR CODE  
■ HALFTONE  
□ NON\_HALFTONE





**Fig. 18D**

COLOR CODE  
■ HALFTONE  
□ NON\_HALFTONE



"we did not do anything that is related to financial structuring or disclosure in any of the issues that got them into trouble. We stand by all the work we did. Beyond that, we can only empathize with the trouble they are going through. It's a sad thing to see."

Still, outsiders marvel that the secretive partnership has not been drawn into the debate, given its extensive involvement at Enron. "I'm surprised that they haven't been subpoenaed as a witness, at least," says Wayne E. Cooper, CEO of Kennedy Information, a research and publishing firm that keeps tabs on consultants. "There was so much smoke coming out of the Andersen smoking gun that all the firefighters went after that one. McKinsey was lucky. They dodged a bullet."

The bad news, however, is that Enron, which was

Did McKinsey's partners get caught up in the euphoria of the late '90s and suffer lapses of judgment? And if so, what does that say about the quality of its expensive advice? Did it stray from its core values? What accountability does it—or any consulting firm—have for the ideas and concepts it launches into a company?

After all, McKinsey was a key architect of the strategic thinking that made Enron a Wall Street darling. In books, articles, and essays, its partners regularly stamped their imprimatur on many of Enron's strategies and practices, helping to position the energy giant as a corporate innovator worthy of emulation. The firm may not be the subject of any investigations, but its close involvement with Enron raises the question of whether McKinsey, like some other professional firms, ignored warning flags in order to keep an important account.

The breakdowns of such visible clients could not have come at a more trying time. Instead of celebrating the end of his third and final three-year term

**1926** University of Chicago accounting professor James "Mac" McKinsey founds the firm. First client: meatpacker Armour & Co.

paying McKinsey as much as \$10 million in annual fees, is just one of an unusual number of embarrassing client failures for the elite consulting firm. Besides Enron, there's Swissair, Kmart, and Global Crossing—all McKinsey clients that have filed for bankruptcy in relatively short order. And those are just the biggest. McKinsey also finds itself im-

All of which raises uncomfortable questions about the world's most prestigious—and enigmatic—consulting firm.

1933 Marvin Bower, a lawyer and Harvard MBA, joins the New York office and essentially invents the discipline of modern management consulting.

1935 Marshall Field & Co. becomes the firm's first big client and eventually hires James McKinsey to be CEO.

1959 Opens first non-U.S. office in London and wins landmark assignment with Bank of England.

1960 Successful expansion overseas throughout Europe and Australia.

1964 Launches *The McKinsey Quarterly*, which remains the foremost journal of management ideas produced by a consulting firm.

1978 McKinsey partner Lou Gerstner departs firm for American Express and later becomes CEO of RJR Nabisco and IBM. Other alumni head some of the world's biggest companies.

as managing director, Gupta, 53, finds his firm roiled by a rare and potentially disruptive downturn in its business. Like most other consulting firms, McKinsey rode the e-business wave to record revenues—and record partner payouts—in 2000. When the boom turned to bust, the firm was stuck with far too many consultants and not nearly enough assignments. The utilization rate, or billable time, of its consultants has fallen to its lowest level in more than 32 years: just 52%, vs. the heady 64% level during the dot-com boom.

That's not to say that McKinsey has lost its standing. The firm remains the high priest of high-level consulting, with the

**McKINSEY OVER THE YEARS**

Fig. 18E



D-ROM※4、彩ちゃんのセッ  
 D-ROMソフト※5 (①カード工  
 吾TureTypeフォント)

異なります。※3:カット紙専用  
 リンタ同梱のシートにてお申し

Fig. 19A

D-ROM※4、彩ちゃんのセッ  
 D-ROMソフト※5 (①カード工  
 吾TureTypeフォント)

異なります。※3:カット紙専用  
 リンタ同梱のシートにてお申し

Fig. 19B

D-ROM※4、彩ちゃんのセッ  
 D-ROMソフト※5 (①カード工  
 吾TureTypeフォント)

異なります。※3:カット紙専用  
 リンタ同梱のシートにてお申し

- COLOR CODE
- FINE EDGE
  - COARSE EDGE
  - HALFTONE
  - CONSTANT TONE
  - OTHER

Fig. 19C





Fig. 20A



Fig. 20B




Fig. 20C

COLOR CODE

- FINE EDGE
- COARSE EDGE
- HALFTONE
- CONSTANT TONE
- OTHER



COLOR CODE

- FINE EDGE
- COARSE EDGE
- HALFTONE
- CONSTANT TONE
- OTHER




Fig. 21A




Fig. 21B




Fig. 21C



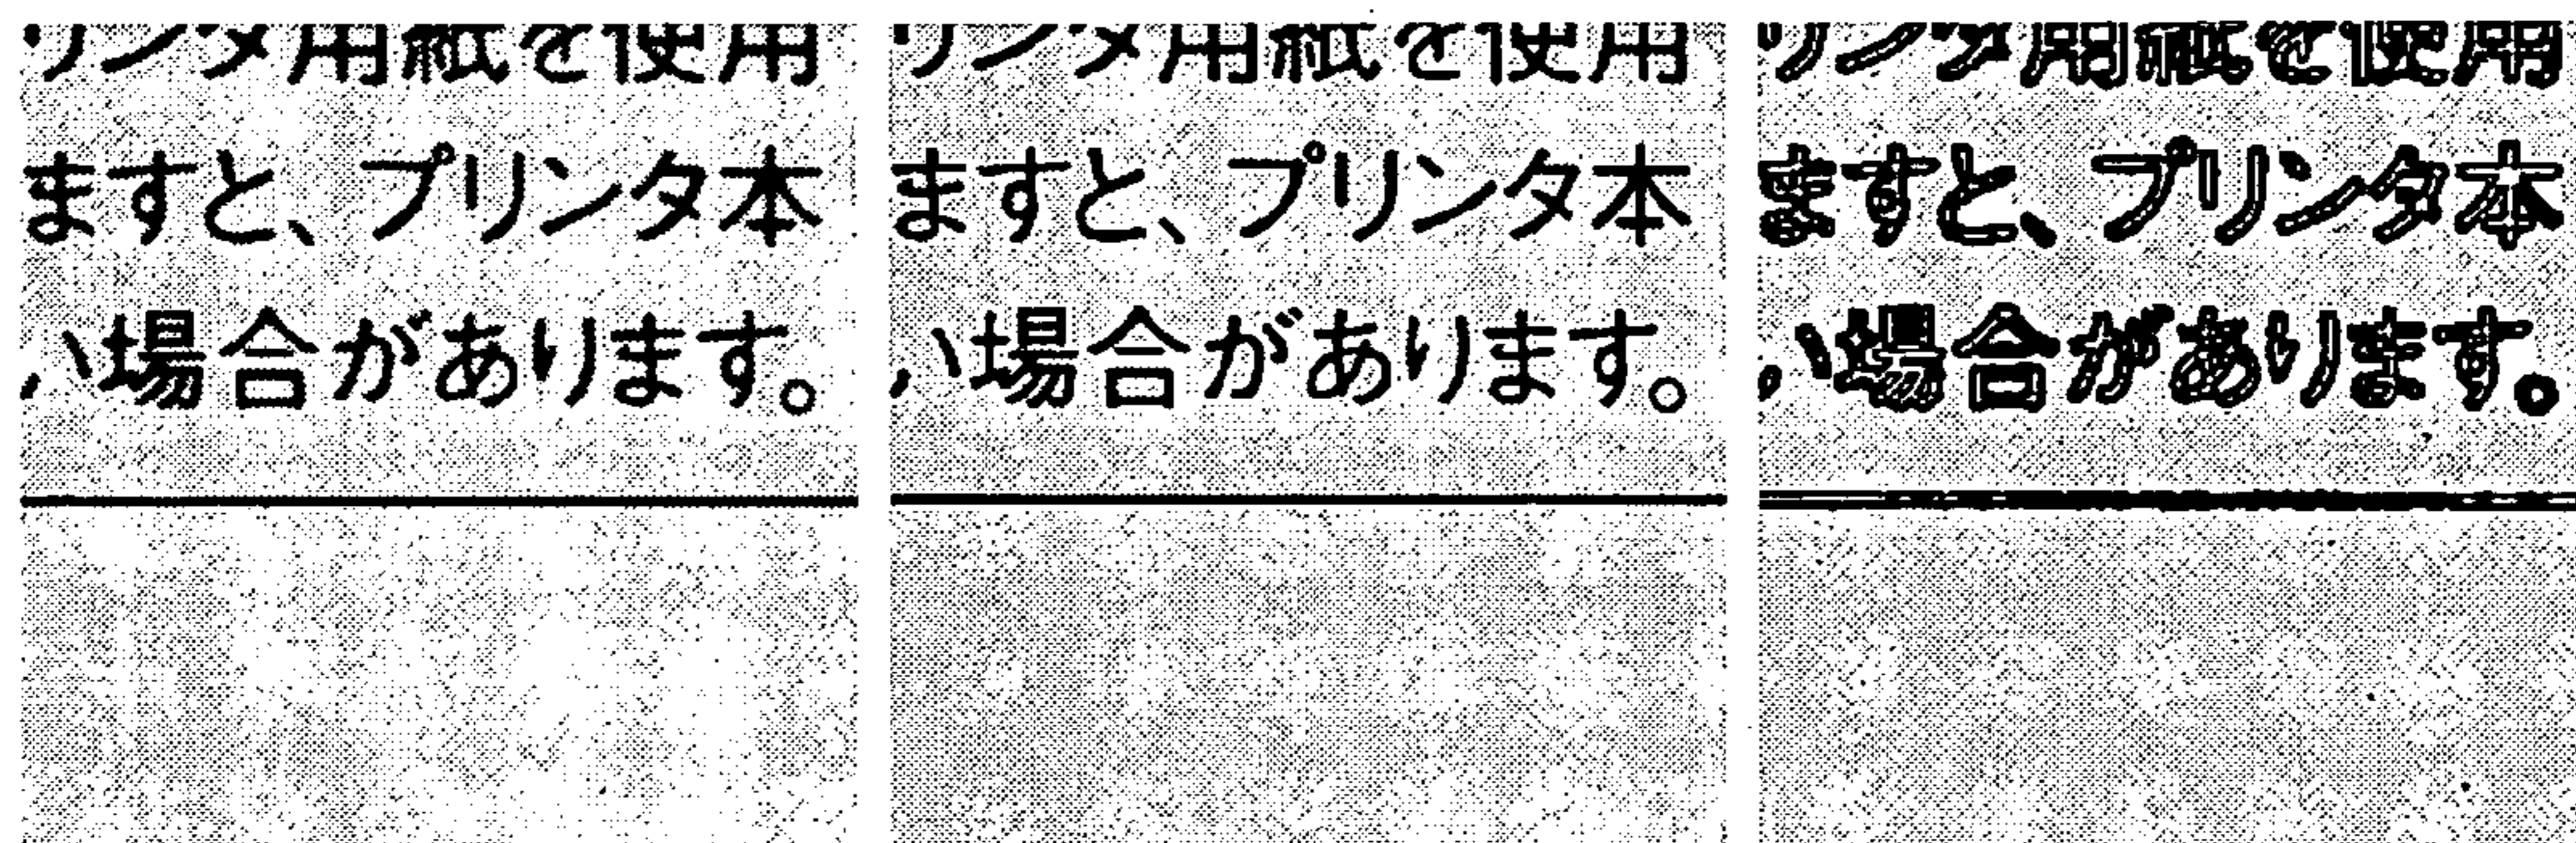


Fig. 22A

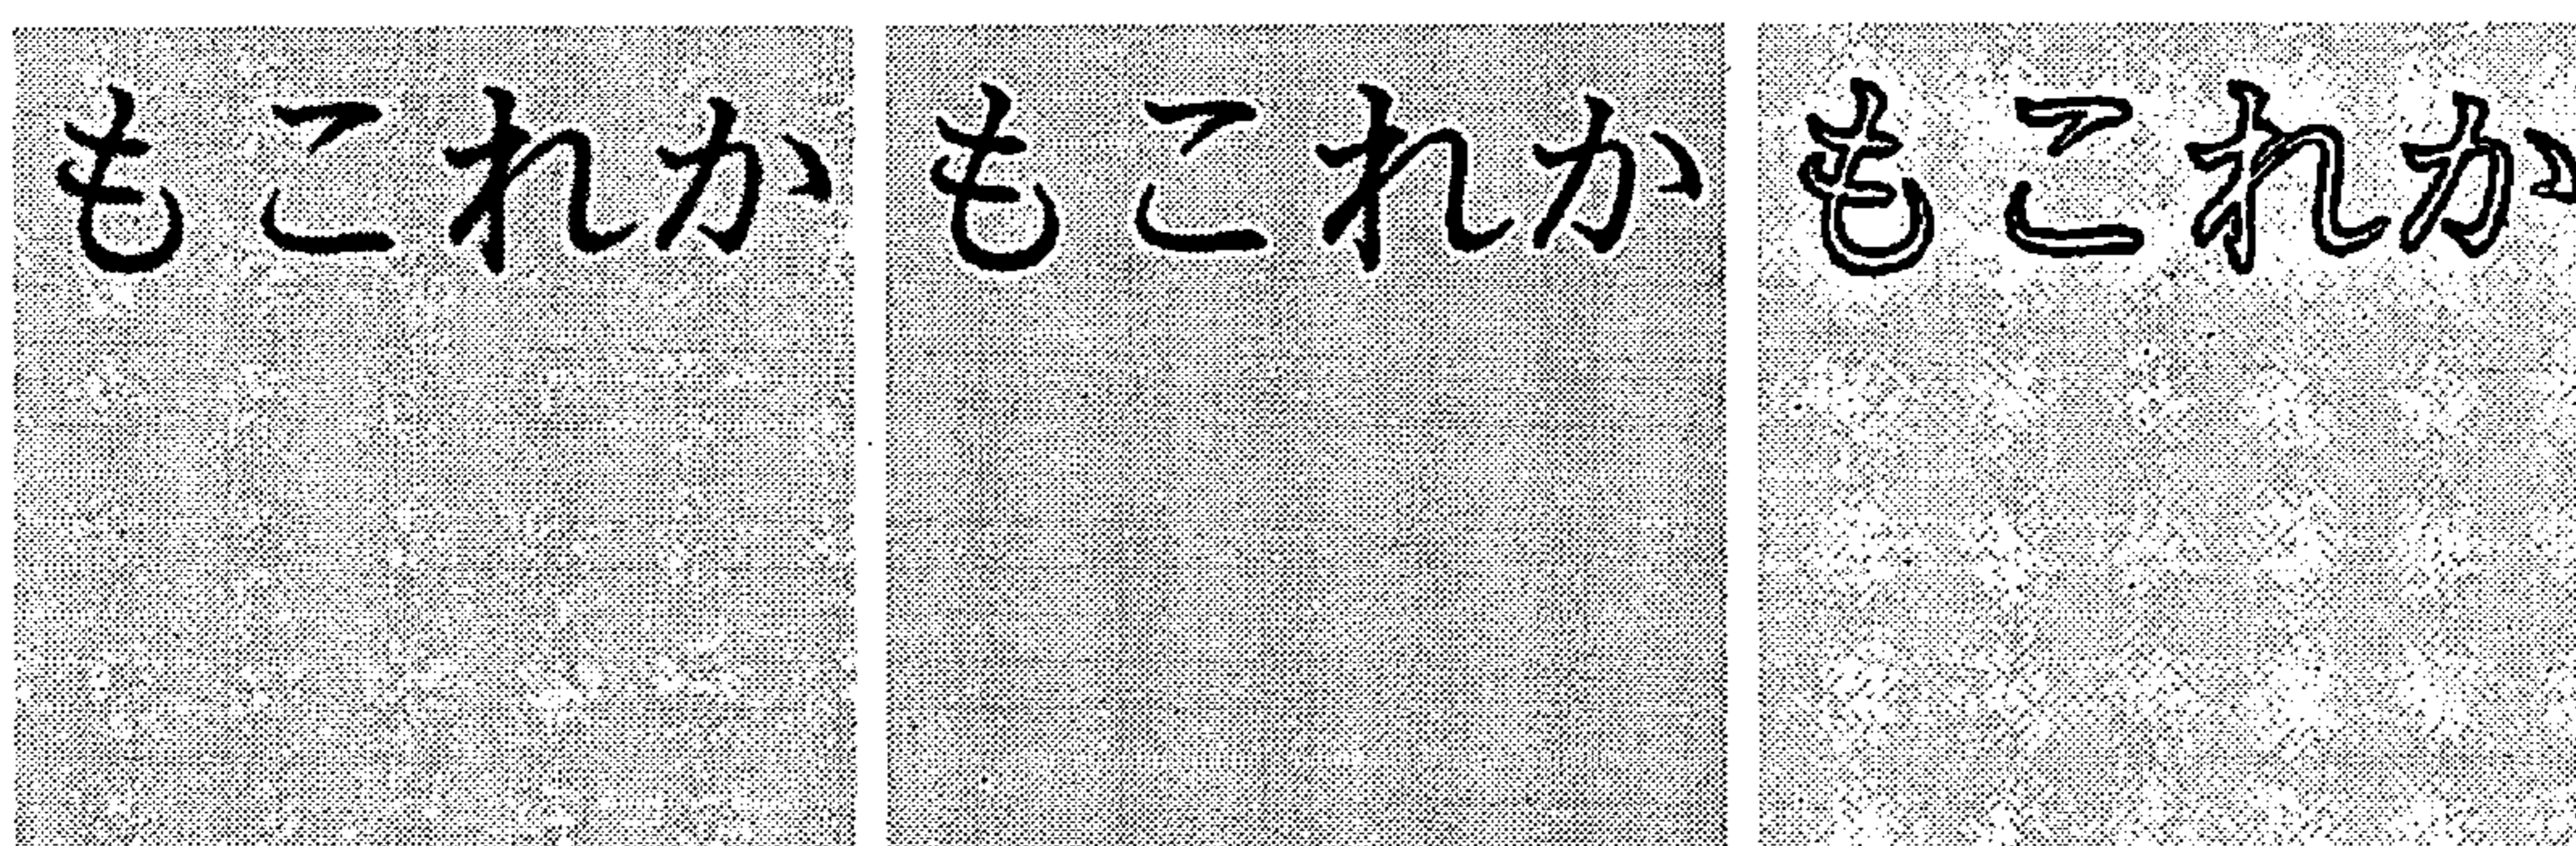


Fig. 22B




Fig. 22C

- COLOR CODE
- FINE EDGE
  - COARSE EDGE
  - HALFTONE
  - CONSTANT TONE
  - OTHER






Fig. 23A




Fig. 23B




Fig. 24



**METHOD AND APPARATUS FOR  
SEGMENTATION OF COMPOUND  
DOCUMENTS HAVING LOW RESOLUTION  
HALFTONES**

CROSS REFERENCE TO RELATED  
APPLICATIONS

This application claims priority under 35 U.S.C. § 119(e) on U.S. Provisional Patent Application No. 60/453,077 filed Mar. 6, 2003, and entitled "LOW COMPLEXITY COMPREHENSIVE LABELING AND ENHANCEMENT ALGORITHM FOR COMPOUND DOCUMENTS," and claims priority under 35 U.S.C. § 120 as a continuation-in-part of U.S. patent application Ser. No. 10/625,107 filed on Jul. 23, 2003 and entitled "METHOD AND APPARATUS FOR SEGMENTATION OF COMPOUND DOCUMENTS." These applications are incorporated by reference herein in their entirety for all purposes.

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates generally to document replication systems and more particularly to a method and apparatus for segmenting and enhancing compound documents for image replication applications.

2. Description of the Related Art

The quest for high quality in everyday copying applications has led to increased sophistication in scanners and printers. In particular, the desire for high quality text and graphics output has forced modern scanners and printers to operate at 600 dots per inch (dpi) or higher dpi resolutions even on basic document applications. However, the associated increases in the costs of these devices leaves little room for the expense of signal processing algorithms which must still accomplish high quality output in an automated mode of operation. The result is a direct need for compound document processing techniques of very low computational and memory complexity.

The constraint for very low cost algorithms is in part alleviated by the high resolution devices that make up a copier. Modern scanners with low noise sensors and printers with precision dot placement no longer necessitate very high performance signal processing for noise removal, deblurring, etc. However, halftone dots in the input become much more pronounced in high resolutions, thereby resulting in possible Moire artifacts that are introduced if such regions are printed unprocessed. FIGS. 1A-1C illustrate various portions of a document scanned at different resolutions. FIG. 1A is scanned at 150 dpi, FIG. 1B is scanned at 300 dpi, and FIG. 1C is scanned at 600 dpi. It should be appreciated that FIGS. 1A and 1B were each scaled to 600 dpi by pixel replication for comparison purposes. As illustrated in successive halftone regions 102a-102c of FIGS. 1A-1C, respectively, the resolution increases from 150 dpi to 600 dpi. While the text and other edges become sharper with the increasing resolution, the halftone dots become more pronounced.

Attempts to segment a compound document to halftone and non-halftone regions have met with limited success. For example, some techniques require pre-segmentation of the document into blocks or other rectangular regions. Additionally, some segmentation techniques look at the document at a single resolution, which inevitably leads to erroneous decisions and may restrict the technique to detecting a particular class of halftones. Furthermore, these techniques

are not applicable to general document types, such as scanned documents, and are associated with a high computational complexity and excessive consumption of memory resources.

5 Additionally, previous attempts mainly concentrate on region classification on documents that have been pre-segmented to their constituent regions. Each one of the pre-segmented regions are classified into various classes and the segmentation is typically delegated to sophisticated and computationally complex document analysis algorithms. Furthermore, irrespective of how simple the region classification is, such approaches demand at least two passes on the data resulting in significant memory and memory-bandwidth requirements. Moreover, the overall operation of pre-segmentation and classification may still require considerable computation. These pre-segmentation based methods range from simple techniques which use an " $\alpha$ -crossing" technique for detecting halftone regions, and techniques which operate on binary only documents and detect halftone regions using predefined masks, to the more elaborate and computationally complex techniques, which utilize Fourier transforms and tuned directional bandpass filters for texture analysis and halftone detection. The previous techniques effectively pit halftone and text regions against one another by basing the identification decision on complex statistics (the alternative hypothesis are text, halftone, non-halftone, etc.) By subtracting most of the influence of real document edges, the segmentation and identification decisions in this work are based on simple statistics (the alternative hypothesis are halftone and non-halftone).

Other work that tries to identify halftones by detecting periodicities is limited to input documents containing certain types of halftones. Moreover, one must be assured that periodicities in the original halftone are sufficiently preserved after the color space change due to the scan of the original printed halftone (for e.g., CMYK to RGB). Techniques that try to detect real document edges that are not due to halftone dots by using edge continuity analysis are severely sensitive to thresholds since disconnected halftone dots may appear connected depending on the threshold used in edge detection.

As a result, there is a need to solve the problems of the prior art to provide a technique that segments a document irrespective of the document characteristics and minimizes the computational cost as well as memory consumption.

SUMMARY OF THE INVENTION

50 Broadly speaking, the present invention fills these needs by defining a multi-resolutional scheme for labeling and decomposing a broad range of documents followed by the selective enhancement of the document, while consuming minimal computational and memory resources. It should be appreciated that the present invention can be implemented in numerous ways, including as a process, a system, or a device. Several inventive embodiments of the present invention are described below.

In one embodiment, a method of labeling digital image data associated with a compound document in a single pass is provided. The method initiates with receiving image data. Then, the image data is distributed to multiple branches, wherein at least two of the multiple branches include filters configured to blur the image data at different resolutions. Next, edges of the blurred image data are detected at the different resolutions. Then, edges of unfiltered image data are detected. Next, data output from the at least two of the



multiple branches are combined. Then, it is determined whether the combined data is to be labeled as a low resolution halftone.

In another embodiment, a method for segmenting a compound document for enhancement during replication of the compound document is provided. The method initiates with filtering data representing a portion of the compound document through a filter associated with a first resolution. Then, the data representing the portion of the compound document is filtered through a filter associated with a second resolution. Next, edges are detected on both, an output of the filter associated with the first resolution and an output of the filter associated with the second resolution. Then, data representing detected edges from both outputs are combined. Next, it is determined whether a pixel corresponding to the data representing the detected edges is over a halftone region.

In yet another embodiment, a computer readable medium having program instructions for segmenting a compound document for enhancement during replication of the compound document is provided. The computer readable medium includes program instructions for filtering data representing a portion of the compound document through a filter associated with a first resolution and program instructions for filtering the data representing the portion of the compound document through a filter associated with a second resolution. Program instructions for detecting edges on both, an output of the filter associated with the first resolution and an output of the filter associated with the second resolution are included. Program instructions for combining data representing detected edges from both outputs and program instructions for determining whether a pixel corresponding to the data representing the detected edges is over a halftone region are provided.

In still yet another embodiment, an image replication system is provided. The image replication system includes a labeling module configured to segment image data corresponding to a compound document. The labeling module includes multiple branches having edge detection modules for detecting edges of the compound document. A plurality of the multiple branches are capable of filtering the image data at different resolutions. The labeling module includes logic for combining output of at least two of the multiple branches in order to initially identify a portion of the compound document as a halftone portion. The labeling module also includes a counting module configured to count pixel values within a neighborhood to determine whether the initially identified halftone portion is finally labeled as a halftone region.

In yet another embodiment, an integrated circuit capable of segmenting image data corresponding to a compound document is provided. The integrated circuit includes labeling circuitry configured to identify an image data region type. The labeling circuitry has multiple branches. Each of the multiple branches includes filter circuitry capable of filtering the image data and edge detection circuitry configured to detect edges of the filtered image data. Each of the multiple branches are associated with a different filter resolution. The labeling circuitry further includes circuitry for combining at least two outputs of the multiple branches and circuitry for identifying the image data region type based upon a value defined through the circuitry for combining the at least two outputs of the multiple branches.

Other aspects and advantages of the invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the invention.

## BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, and like reference numerals designate like structural elements.

FIGS. 1A–1C illustrate various portions of a document scanned at different resolutions.

FIG. 2 is a simplified schematic diagram illustrating a copier pipeline configured to execute the segmentation, labeling and selective enhancement functionality in accordance with one embodiment of the invention.

FIG. 3 is a simplified schematic diagram illustrating the banded nature of the technique to minimize memory resources in accordance with one embodiment of the invention.

FIG. 4A is a simplified schematic diagram of the labeling process with exemplary images provided for illustrating the effects of the processing in accordance with one embodiment of the invention.

FIGS. 4B-1 through 4B-4 illustratively represent the logical AND operation and the subtracting operation discussed above with reference to FIG. 4A.

FIG. 4C is a simplified schematic diagram of an alternative embodiment of the labeling process of FIG. 4A in accordance with one embodiment of the invention.

FIGS. 5A and 5B illustrate an original image portion and a low pass filtered image portion of the original, respectively.

FIGS. 6A and 6B illustrate the low pass filtered image of FIG. 5B and detected coarse scale edges from the low pass filtered image, respectively, in accordance with one embodiment of the invention.

FIGS. 7A–7E are schematic diagrams illustrating an edge proximity technique for detecting missed edges during coarse edge detection on a schematic text character in accordance with one embodiment of the invention.

FIGS. 8A and 8B illustrate an original scanned image portion and a corresponding labeled fine scale image portion, respectively, in accordance with one embodiment of the invention.

FIG. 9 illustrates an image portion having locations marked as fine edges but not as coarse edges for the example document in accordance with one embodiment of the invention.

FIG. 10 is a schematic diagram illustrating the counting neighborhood utilized for halftone decisions in accordance with one embodiment of the invention.

FIGS. 11A–11D represent an original image scanned at 600 dpi, labeled coarse scale edges, labeled fine scale edges, and fine but not coarse edges, respectively.

FIG. 12 is a schematic diagram illustrating the partial counting neighborhoods in halftone decisions in accordance with one embodiment of the invention.

FIGS. 13A and 13B are schematic diagrams illustrating the per-pixel labels obtained for the examples depicted in FIG. 5A and FIG. 11A, respectively.

FIGS. 14A and 14B illustrate the original scanned document and the enhanced output, respectively, in accordance with one embodiment of the invention.

FIGS. 15A and 15B include a series of images illustrating the application of the embodiments described above for detecting different types of halftones in accordance with one embodiment of the invention.

FIGS. 16A–16E illustrate performance of the embodiments described herein on a multilingual document containing halftones and fine text.



## 5

FIGS. 17A–17E illustrate performance of the embodiments described herein on a multilingual document having a very irregularly shaped halftone region.

FIGS. 18A–18E again shows performance of the embodiments described herein on a document with an irregularly shaped halftone region.

FIGS. 19A–19C show the enhancement results under situations due to Case 1 and Case 3 type errors in accordance with one embodiment of the invention.

FIGS. 20A–20C illustrate results over Case 4 situations in accordance with one embodiment of the invention.

FIGS. 21A–21C show the performance of the technique over Case 5 problems as seen with reference to FIG. 18B.

FIGS. 22A–22C are a schematic diagrams illustrating the technique's performance on bleed-through removal.

FIG. 23A is a simplified schematic diagram illustrating a portion of a compound document which may be mistakenly characterized as a halftone region.

FIG. 23B is a simplified schematic diagram illustrating a scheme to address the mistaken characterization discussed with reference to FIG. 23A, through edge detection at three scales in accordance with one embodiment of the invention.

FIG. 24 is an alternative representation of the effect of an additional branch being included in the labeling module as described with reference to FIG. 4C, in accordance with one embodiment of the invention.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

An invention is described for an apparatus and method for providing a stable output while minimizing moire patterns and visible structures when printing images from a printing device. It will be apparent, however, to one skilled in the art, in light of the this disclosure, that the present invention may be practiced without some or all of these specific details. In other instances, well known process operations have not been described in detail in order not to unnecessarily obscure the present invention. FIGS. 1A–1C are described in the “Background of the Invention” section.

The embodiments of the present invention provide an apparatus and method for segmenting, labeling and then enhancing an image consisting of various regions of pixels, e.g., text/graphics, images halftones, text/graphics on halftones, constant tones, etc. The image is decomposed into various types of regions so that enhancement techniques optimized for a particular region may be applied to that region. It should be appreciated that the enhanced image may be halftoned and printed, displayed on a monitor, compressed and stored, etc. That is, the results of embodiments described herein may be sent to a multitude of devices and converted into a multitude of formats. The scheme enabling the embodiments described is multi-resolutional in nature, i.e., the scheme looks at the document in more than one resolution to more accurately segment the document. In one embodiment, the scheme defined herein performs processing on the input (fine scale) and multiple low pass filtered versions of the input (coarse scale). Edge detection is done at multiple resolutions to mark “coarse edges” on the low pass filtered image and with non-low pass filtered data to mark “fine edges” on the original image. With the multiple low pass filters chosen to suppress most of the halftone, the edge detection on the low pass filtered image marks most of the real edges (due to image edges, text, graphics, etc.), while fine edge detection marks the real edges and halftone dots. All locations that are-marked as fine edges but not as coarse edges serve as initial estimates of

## 6

halftone dots. Such locations are counted around each pixel and the result thresholded in order to determine halftone pixels. Coarse edges on detected halftone regions and fine edges on non-halftone regions make up two of the labels in the X field. In another embodiment, a variance is calculated around each non-halftone pixel to determine if that pixel is over a constant-tone region. As a result of this process we have the five possible labels in Table 1.

Once labeling is done, edge enhancement is carried out on the corresponding detected coarse and fine edge pixels and descreening is accomplished by low pass filtering the halftone pixels. Descreening involves low pass filtering, therefore, the scanned document is segmented or decomposed to its constituent halftone and non-halftone regions in order not to corrupt high resolution text and graphics with low pass filtering. The embodiments described below accomplish this segmentation in conjunction with error tolerant descreening and enhancement in a very low complexity environment. In one embodiment, bleed-through removal is applied to the enhanced data. As will be explained in more detail below computations may be reused in the enhancement stage, which in turn share some of the computational results of the labeling stage.

FIG. 2 is a simplified schematic diagram illustrating a copier pipeline configured to execute the segmentation, labeling and selective enhancement functionality in accordance with one embodiment of the invention. Image replication device 108, e.g., a copier, includes scanner module 110, integrated enhancement module 112, and printing module 120. The data coming in from scanner module 110 is labeled, descreened, and enhanced through labeling module 114, descreening module 116, and text/graphics enhancement module 118, respectively. The processed data is output together with the detected per-pixel labels to printing module 120. It should be appreciated that the labels comprise an extra channel in addition to the RGB channels in the output stream that is referred to as the X channel or X field. In one embodiment, the X channel is configured to aid further operations in the printing pipeline that may benefit from the labeling, e.g., the adaptive determination of printer halftone screens, resolution improvement operations, etc. Of course, the RGB data stream is exemplary and the functionality described herein may operate on other types of color documents, as well as on grayscale documents and on black and white input. Table 1 illustrates the exemplary per-pixel labels (X field) detected at the labeling stage of the algorithm and summarizes the labels used for the X field in one embodiment of the invention. It should be appreciated that the detected edges are those edges determined to be due to image edges, text, graphics, etc. That is, the detected edges are not due to the halftone dots.

TABLE 1

55 FINE EDGE	A detected edge pixel over a non halftone region in the document.
COARSE EDGE	A detected edge pixel over a halftone region in the document.
HALFTONE	A detected (non-edge) halftone pixel in the document.
60 CONSTANT TONE	A detected constant-tone pixel over a non halftone region in the document.
OTHER	Marks all the remaining pixels in the document.

It should be further appreciated that the technique described herein, also referred to as an algorithm, makes no attempt at distinguishing among edges due to text, graphics, images, etc. While the primary application of the function-



ality is the accomplishment of the labeling and enhancement for a copier pipeline, the results of the algorithm may be used in conjunction with more sophisticated algorithms, e.g., optical character recognition, etc., to enable further document processing and segmentation. Similarly, the labeling results may be used to derive other applications such as the compression of compound documents, etc.

It will be apparent to one skilled in the art that there are many other operations that may be performed on the output stream as it progresses in the printing pipeline of printing module **120**. For example, in some modular copying applications, where the scanner and printer are well separated, it may be desirable that the deployment of the labeling, descreening, and enhancement algorithm not require any changes in the printing pipeline. On the other hand in applications where scanner-side and printer-side algorithms are jointly designed, the technique must be flexible enough not to curtail the possible savings introduced by the joint design. This flexibility is provided through the embodiments described herein. For example, the technique may operate on raw scanner output without requiring extra color correction, thereby allowing separate scanner and printer color correction to be combined for computational savings. Furthermore, the output X channel may be suppressed in modular applications. The technique may thus be incorporated in the scanner-side or the printer-side in order to give a system designer freedom to best take advantage of cost, complexity, and performance tradeoffs.

As is generally known, printers use screens to generate an image to be printed. Therefore, if text regions and graphics/image regions are identified, the printer may apply a lower frequency screen to an image region in order to get better color production. On the other hand, the text regions may use high frequency screens in order to make the text appear sharp. Thus, by providing the optional X field to a printer module, the printer module could then execute functionality on top of the enhanced image data provided by the embodiments described herein to further enhance the image for printing.

FIG. **3** is a simplified schematic diagram illustrating the banded nature of the technique to minimize memory resources in accordance with one embodiment of the invention. Operating in a banded fashion refers to having a relatively small number of rows stored in memory and all processing is done on the stored rows. As logically represented in FIG. **3**, band of stored rows **122** in memory **130** moves over the input, creating one row of output for each row of input. However, it should be appreciated that the algorithm is not limited to banded processing as any suitable processing scheme may be utilized. Shaded region **124** inside band of rows **122** denotes the active region used in processing which slides from left to right within band of stored rows **122** as the columns in the band are processed. In one embodiment, at any point in the processing loop only information that is within active region **124** is utilized. The limited extent of active region **124** assists in curbing the active memory requirements associated with the technique and also allows integration within systems that perform block-based rather than band-based processing. Depending on the various parameters, the number of stored rows may change. For the 600 dpi examples in this paper, the algorithm is configured to store 17 rows and the active region corresponds to approximately  $17 \times 257 = 4,369$  pixels. Compared to full page memory at 600 dpi given by  $8.5 \times 11 \times 600 \times 600 = 33,660,000$  pixels, in banded operation the algorithm requires storage on three buffers over  $17 \times 8.5 \times 600 = 86,700$  pixels which is about 0.8% of full page memory. The

active region used by the algorithm is about 0.01% of full page memory. Thus, the thickness of the rows stored by the algorithm (i.e., the size of the band in row dimension) is comparable to the size of a full-stop in 11 pt or 12 pt text. It should be appreciated that the embodiments described herein are not limited to a banding scheme where 17 rows are used. When banded processing is used any suitable number of rows may be employed.

It should be appreciated that the technique described herein is configured to perform a single pass on the input and once the band reaches the bottom of the input page, the output page (which includes the enhanced input and the optional X field) is ready up to a few rows of delay. Accordingly, the operation of this algorithm may be analogized to a very thin line that rapidly scans an input page while simultaneously producing an output page. The memory requirements and memory-bandwidth requirements of the algorithm are minimal due to the banded operation with small active memory. This architecture, in conjunction with the utilized low complexity algorithms, allows very fast processing at a very low cost.

In one embodiment, the algorithm maintains three main buffers that store information corresponding to pixels that are within the band of FIG. **3**. These three buffers store relevant portions of the input image, the low pass filtered version of these portions, and per-pixel labels that have been assigned to these portions. Each stage of the algorithm is described in more detail below with exemplary results of the corresponding stage illustrated on an example document. Color channels in the input document are processed in turn. The primary target of the algorithm is 600 dpi processing, however, in order to demonstrate the flexibility of the algorithm, the various parameter values may be adjusted to other suitable dpi resolutions. It should be appreciated that there may be places where the algorithm makes labeling mistakes. While many of these mistakes may be corrected using straightforward processing, such as median filtering, region completion, etc., it is important to note that the utilized enhancement methods are robust to these errors as described in more detail below, thereby making special processing for these corner cases redundant.

The technique is designed around the observation that the main information in documents is meant to be visible to human observers viewing the documents from a certain distance. Accordingly, the coarse resolution image should contain the bulk of this information and it can be used to identify the information bearing portions at the fine resolution. Therefore, most of the influence of real document edges may be subtracted, through the coarse resolution, and segmentation and identification decisions may be based on relatively simple statistics where the alternative hypothesis are halftone and non-halftone. In contrast, traditional techniques have to base statistical decisions on complex statistics where they have to distinguish among real document edges, non-halftone areas, and halftone areas. Depending on the language, characters, fonts, etc., text and halftone regions can have remarkably similar statistics which limits the robustness of these traditional techniques. Consequently, the algorithm is much more robust and it is applicable over a variety of text and graphics regions, since most of the visual information that is meant to be seen in the input document never enters the statistical decision process to cause mistakes. Similarly, the technique discussed herein incorporates robustness in terms of the types of halftones it may detect since, the halftones are primarily located as regions that have densely packed but invisible information.



FIG. 4A is a simplified schematic diagram of the labeling process with exemplary images provided for illustrating the effects of the processing in accordance with embodiments of the invention. Scanned image document 132, a portion of which is represented by image 132a, is provided to processing branches 133 and 135. Processing branch 133 processes scanned image 132 through low pass filter 134. The output of low pass filter 134 is input into edge detector 136 which is configured to detect coarse edges. It should be appreciated that low pass filter 134 yields a blurred document as illustrated by image portion 134a. In essence, low pass filter 134 suppresses halftones, since high frequency data is filtered out by the low pass filter. In other words, low pass filter 134 functions similar to a human eye in that it detects only low frequency data. Consequently, edge detector 136 detects coarse edges that a human eye would detect. In one embodiment, text, lines, and text on halftone is picked up by edge detector 136. A portion of the image output from edge detector 136 is illustrated by image portion 136a.

Processing branch 135 of FIG. 4A differs from processing branch 133 in that processing branch 135 does not include a low pass filter. Consequently there is no suppression of halftones through processing branch 135. Here, text as well as halftones are picked up through edge detector 138, illustrated by image portion 138a. Thus, processing branch 133 performs the processing at a lower frequency resolution, while processing branch 135 performs the processing at the original resolution of scanned document 132. It should be appreciated that if only processing branch 135 was used without processing branch 133 then the uniform enhancement of the output from edge detector 138 would enhance the halftones as well as detected edges.

Still referring to FIG. 4A, processing branch 133 outputs low pass filtered image data that has coarse edges detected, represented as image A in box 137. Processing branch 135 outputs image data that has fine edges detected, represented as image B in box 139. The data representing image A 137 and image B 139 are combined through logic module 140. In one embodiment, the data is combined through a simple AND operation. It should be appreciated that the combination of the data representing image A 137 and image B 139 yield image data in which the real edges are highlighted. Logic module 140 takes the ANDed output of image A 137 and image B 139 and subtracts the ANDed output from image B 139 to yield data representing image C 141. It should be appreciated that whenever there is a real text edge in image portion 138a, the corresponding location in image portion 136a also is a real text edge. However, where there is a halftone edge in image portion 138a, there will be no corresponding edge in image portion 136a. Therefore, through the AND operation where there is a "1" in corresponding locations of image portions 136a and 138a, the AND output is "1." Where there is a "0" in either of image portions 136a and 138a, then the AND output is "0". Therefore, the real edges are highlighted because halftones are not visible in image portion 136a. In turn, when the ANDed output of the data representing image portions 136a and 138a are subtracted from the data representing image portion 138a, the result corresponds to all the locations that are halftone regions.

FIGS. 4B-1 through 4B-4 illustratively represent the AND operation and the subtracting operation discussed above with reference to FIG. 4A. FIG. 4B-1 corresponds to image portion 138a of FIG. 4A. Here, the text (T) is over halftone region 164-1. Region 164-1 includes halftone dots 162 as the halftones are not suppressed, while region 166-1 corresponds to a continuous tone region. The halftone dots are

associated with a logical "1," while remaining points are associated with a logical "0." FIG. 4B-2 corresponds to image portion 136a in that the low pass filter has been applied to suppress halftones, therefore, there are no halftone dots as in FIG. 4B-1. However, the text images, i.e., "T" and "Z," of halftone region 164-2 and continuous tone region 166-2, respectively, are somewhat blurred due to the low pass filter operation.

FIG. 4B-3 represents the result of the AND operation between the data representing FIGS. 4B-1 and 4B-2. One skilled in the art will appreciate that the AND operation replicates the "T" in halftone region 164-3 and the "Z" in continuous tone region 166-3. That is wherever a logical "1" exists in corresponding regions of FIGS. 4B-1 and 4B-2, an edge is defined is FIG. 4B-3. It should be appreciated that the text images are sharper than in FIG. 4B-2, but not quite as sharp as in FIG. 4B-1, and that the halftone dots from FIG. 4B-1 have been eliminated. FIG. 4B-4 represents the subtraction of the data from the AND operation, i.e., FIG. 4B-3, from the data where the halftones have not been suppressed, i.e., FIG. 4B-1. It will be apparent to one skilled in the art that a dot will be defined through the subtraction operation where there is a "1" in the data image of FIG. 4B-1 and a "0" in the data image representing the AND operation, i.e., FIG. 4B-3. Thus, the operation described with reference to FIGS. 4B-1 through 4B-4 represent a method for effectively selecting halftone dots. It should be appreciated that there may be some mistakes in picking up the real edges, e.g., dots 170 of FIG. 4B-4. However, the remaining functionality discussed below addresses the correction of these errors.

Returning to FIG. 4A, the image data represented by image C 141, e.g., image portion 140a, is communicated to counting module 142, where a counting algorithm is executed. The counting algorithm establishes a neighborhood at each pixel associated with the image data of image C 141. For example, region 172 of FIG. 4B-4 represents the neighborhood around the pixel corresponding to dot 170. The counting algorithm will then go and count the ones in that neighborhood. If the value associated with the number of counted ones is greater than a threshold of threshold module 144, then the corresponding pixel is in a halftone region. If the value associated with the number of counted ones is less than a threshold value, then the corresponding pixel is not in a halftone region. The threshold may be determined experimentally, e.g., through running experiments to determine an optimum threshold. In another embodiment, the threshold is effectively the density of halftone dots. Therefore, where the resolution is known, i.e., the number of halftone dots in a neighborhood may be determined, the threshold may be calculated based on this determination. It should be appreciated that the counter is configured to correct any mistakes through the neighborhood application as discussed earlier.

The output of threshold module 144 of FIG. 4A is communicated to either branch 145 or branch 147. If a pixel is associated with a halftone region, then the pixel will either be labeled as a coarse edge, i.e. an edge on a halftone, or just a halftone in branch 147. In one embodiment, the determination of whether the pixel is a coarse edge or a halftone is made by checking the output of edge detector 136. That is, if edge detector 136 picked up a logical "1" for a point associated with that pixel, then that pixel is a coarse edge. Otherwise, the pixel is labeled as a halftone. Regarding branch 145, if a pixel is associated with a non halftone region, the determination of whether the corresponding pixel is labeled as an edge is made by checking the output of edge



detector **138**. If the corresponding pixel is labeled as an edge in the output of edge detector **138**, then the pixel is labeled as a fine edge, an edge in a non-half-tone region. If the corresponding pixel is not labeled as an edge in the output of edge detector **138**, then variance calculation module **148** is utilized to define a label for the corresponding pixel.

Variance calculation module **148** of FIG. **4A** is configured to determine whether non-half-tone regions are disposed over slowly varying portions of the document, i.e., basically constant regions of the document such as constant color regions. It should be appreciated that bleed-through, the incorporation of shading from the backside of a document interfering with the copying of a front side of the document, is substantially eliminated through the application of variance calculation module **148**. Here again, a neighborhood is established around the corresponding pixel, then the mean and variance are calculated for the pixels in that neighborhood. If the variance is below a threshold variance, then the pixels in the neighborhood are slowly varying, i.e., they are close in value together. Thus, the pixel is labeled as a constant tone pixel. If the variance is greater than a threshold variance, then the pixel is labeled as an other pixel. That "other" pixel label will be described in more detail below. It should be appreciated that box **146** represents the various labels for the X field in one embodiment of the invention. In addition, the labels are meant to be illustrative and not limiting.

It will be apparent to one skilled in the art that the functionality described with reference to box **154** of FIG. **4A** may be incorporated into a microprocessor, i.e., embodied as hardware for an image replication device. For example, the functionality may be provided through a plug-in board, or the like, having a microprocessor configured to perform the above described functionality. Alternatively, the functionality may be provided in the form of software code embodied in a driver that may be downloaded to an image replication device over a distributed network, e.g., the Internet. Thus, the embodiments described herein may be captured in any suitable form or format that accomplishes the functionality described herein and is not limited to a particular form or format.

FIG. **4C** is a simplified schematic diagram of the labeling process of FIG. **4A** in accordance with preferred embodiments of the invention. Scanned image document **132**, a portion of which is represented by image **132a**, is provided to processing branches **133'**, **133** and **135**. Processing branch **133'** processes scanned image **132** through low pass filter A **134'**, which is relatively strong and therefore passes only low frequency data; halftones are suppressed. The output of low pass filter **134'** is input into edge detector **136'** which is configured to detect major or coarse edges. The output of edge detector **136'** is represented as image A' **137'** which contains only low frequency data without halftones. Processing branch **133** is configured to enable identification of low-resolution halftones, which may be more common with newspaper articles and images. Here, the output of low pass filter B **134** is delivered to edge detector **136**. Low pass filter B **134** has a higher cut-off frequency, and hence is weaker, than low pass filter A **134'**. Thus, the output of edge detector **136**, represented as image A **137**, includes somewhat higher frequency data in which low resolution halftones are detectable.

Processing branch **135** in FIG. **4C** differs from processing branches **133'** and **133** in that processing branch **135** does not include a low pass filter. Consequently, there is no suppression of halftones through processing branch **135**. Here, text as well as low and high resolution halftones are

picked up through edge detector **138**, illustrated by image portion **138a** and as discussed above with reference to FIG. **4A**.

Logically subtracting the ANDed result of image data A' **137'** and image data A **137** from image data A **137** in logic module **140'** yields edge-detected data with low resolution halftones. It should be appreciated that this result, represented by image C' **141'**, yields image data from which low resolution halftones are capable of being labeled as such, as described in more detail with reference to FIGS. **23A** through **24**. Logically subtracting the ANDed result of image data A **137** and image data B **139** from image data B **139** in logic module **140** yields image C **141'**, which includes edge-detected data with high resolution halftones.

Thus, by including three processing branches, halftones are located at two different resolutions: a higher resolution as is the case with the embodiments described with respect to FIG. **4A**, and a lower resolution as further enabled by the embodiments of FIG. **4C**. Separating the halftone detection process into two makes the process more robust; the halftone detection scheme is optimized for a broader range of halftone frequencies. Moreover, by enabling the detection of low-resolution halftone regions, weak enhancement (also referred to as resolution dependent enhancement/blurring) may be subsequently applied as illustrated with reference to FIGS. **23B** and **24**. The embodiments of FIG. **4C** enable easier identification and hence enhancement of transition areas between halftone and non-half-tone regions, so that artifacts between transitions from halftone to non-half-tone regions on the compound document can be smoothed.

It will be apparent to one skilled in the art that any suitable number of additional branches may be included in the labeling schemes of FIGS. **4A** and **4C**. The additional branches will, of course, come along with increased complexity, however, the greater the number of branches, the enhancement may be applied with more refined control.

Still referring to FIG. **4C**, the image data represented by C' **141'** is communicated to counting module **142'**, where a counting algorithm is executed as discussed above with reference to counting module **142** of FIG. **4A**. That is, the counting algorithm establishes a neighborhood at each pixel associated with the image data C' **141'**. In one embodiment, the counting algorithm counts a logical value, e.g., the ones in the neighborhood. If the value associated with the number of counted ones is greater than or equal to a threshold of threshold module **144'**, then the corresponding pixel is in a halftone region, i.e., a low resolution halftone region. If the value associated with the number of counted ones is less than a threshold value, then the corresponding pixel is not in a low resolution halftone region. The threshold may be determined as described above with reference to FIG. **4A**.

The output of threshold module **144'** is communicated to module **146'**. It should be appreciated that module **146'** represents the various labels for the X field in one embodiment of the invention. It should be further appreciated that the labels are meant to be illustrative and not limiting here.

One skilled in the art will appreciate that the functionality described with reference to module **154'** of FIG. **4C** may be incorporated into a microprocessor or integrated circuit, i.e., an embodiment as hardware for an image replication device. For example, the functionality may be provided through a plug-in board, or the like, having a circuitry configured to perform the above-described functionality. It will be apparent to one skilled in the art that the functionality described herein may be synthesized into firmware through a suitable Hardware description language (HDL). Alternatively, the functionality may be provided in the form of software code



embodied in a driver that may be downloaded to an image replication device over a distributed network, e.g., the internet. Thus, the embodiments described herein may be captured in any suitable form or format that accomplishes the functionality described herein and is not limited to a particular form or format.

The two stages of the algorithm, i.e., labeling and enhancement/descrining are discussed below. The algorithm is primarily based around the labeling stage with the enhancement stage strongly coupled to labeling for added robustness and computational simplicity. As used below let  $S(i,j)$ ,  $i=0, \dots, N-1, j=0, \dots, M-1$  denote a color channel in the original input.

The algorithm uses a set of internal labels that are assigned to pixels through independent classification stages that operate on each color channel of the input. A buffer maintains the results of the independent classifications and these results are combined at the final decision stage to assign a final label for all the color channels in each pixel. It should be appreciated that until the final decision stage, a pixel may have several attached labels for each color channel. The input is initially low pass filtered to yield a coarse resolution image as illustrated in FIGS. 5A and 5B via the following equation:

$$L(i,j)=G(i,j)*S(i,j) \quad (1)$$

where \* denotes convolution and  $G(i,j)$  is the utilized low pass filter. As mentioned above, the algorithm performs edge detection on the low pass filtered data to determine the locations of the real document edges. Hence, it becomes important for the low pass filter to suppress the halftone dots that are present in the document to the degree that these dots are not falsely detected during coarse edge detection. Another issue in the choice of the filter concerns the alignment of detected edges. Since edge detection at coarse and fine scales are used in halftone detection, the detected edges at the two scales should align in order to avoid complex logic in the implementation.

In one embodiment, the alignment requirement is met by using symmetric, zero-delay filters. This also reduces the number of multiplications that are required to calculate the filtered value. The suppression of the halftone dots is not as constraining on the design and the robustness of the design. Halftones may be detected in the input to the degree that they are likely to cause Moire artifacts in the final printed output. Hence, halftone-like features below a certain frequency need not be detected and therefore need not be suppressed by the filter. It should be appreciated that the exact frequency at which this occurs is device specific and depends on the printer used in the copier system, the utilized halftoning method, etc. Furthermore, the halftones present in the input document are designed to be invisible to observers viewing the document from a certain distance. These factors allow straightforward decisions on the choice of the bandwidth of the filter based on the dpi resolution of operation. Very good performance has been experienced using a symmetric Gaussian filter with the bandwidth chosen at 600 dpi and where this bandwidth is adjusted linearly for other target dpi resolutions. For example if the standard deviation of the Gaussian filter is set to  $\sigma_g=2$  at 600 dpi, then the standard deviation of the filter at 300 dpi is set to 1, etc. Note that since a copier device is expected to operate at a fixed set of resolutions, one can also utilize different filters for each resolution with the filter bandwidths adjusted using a train-

ing set at each resolution. All the 600 dpi illustrated herein utilize a symmetric, separable,  $7 \times 7$  Gaussian filter of standard deviation  $\sigma_g=2$ .

After low pass filtering, coarse edge detection is carried out on the low pass filtered image using straightforward edge detection as illustrated in FIGS. 6A and 6B. FIG. 6A illustrates the low pass filtered image of FIG. 5B which corresponds to the output of operation 134 with reference to FIG. 4A. FIG. 6B illustrates detected coarse scale edges of the low pass filtered images and corresponds to the output of operation 136 with reference to FIG. 4A. In order to combat the rare occurrence of unsuppressed halftone dots being detected as edges, the detected edge labels are processed to remove mistakes that manifest themselves as an isolated pixel labeled as an edge. This step is typically only required if the input is expected to contain very low frequency halftones but in a limited fashion it also serves to reduce the dependency on the threshold used in the edge detection. For example, the multiple branches discussed with reference to FIG. 4C, addresses low frequency halftone detection issues.

The edge detection threshold may be adjusted experimentally so that most of the real edges in documents are detected. In one embodiment, the threshold is dependent on the color contrast characteristics of the scanner and whether color correction has been done on the scanned data. The examples used herein are on raw scanner output with no color correction. For the examples we have utilized the directional edge filters given by:

$$E_h(i, j) = \begin{bmatrix} .25 \\ .50 \\ .25 \end{bmatrix} \times \begin{bmatrix} -1 & 1 \end{bmatrix}, E_v(i, j) = \begin{bmatrix} -1 \\ 1 \end{bmatrix} \times \begin{bmatrix} .25 & .50 & .25 \end{bmatrix} \quad (2)$$

for horizontal and vertical edge detection, respectively. The absolute values of results of filtering with these filters was thresholded using the threshold  $T_{thick}=3$  at 600 dpi (300 dpi processing uses the same threshold). An edge is declared if the absolute value output of either of the filters passes the threshold, i.e., given by:

$$D_{c,h}(i,j)=E_h(i,j)*L(i,j), D_{c,v}(i,j)=E_v(i,j)*L(i,j) \quad (3)$$

there is a coarse edge at pixel  $(i,j)$  in the representative channel if  $A(i,j)=1$  where

$$A(i, j) = \begin{cases} 1 & |D_{c,h}(i, j)| \geq T_{thick} \text{ or } |D_{c,v}(i, j)| \geq T_{thick} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

One of the key problems with traditional techniques in the enhancement of edges over halftones is the thickness of the detected edges. Often times, it may be desirable to restrict enhancement to "thinned" edges whereas halftone decisions are still made with the "thick" labels. In order to accommodate such enhancement concerns two thresholds are utilized and distinguished between thick and thin coarse edges using two labels, i.e., a pixel that passes both thresholds is marked as thick and thin. Thick edges are utilized in halftone decisions whereas, as described below, a selective combination of thick and thin edges are utilized in enhancement. The larger threshold used to mark thin edges is given by  $T_{thin}=7$  at 600 dpi ( $T_{thin}=11$  at 300 dpi).

As pointed out above there is a need to distinguish between thick and thin coarse edges in order to accommo-



date enhancement concerns. While halftone detection proceeds by utilizing thick coarse edges, i.e., the determination of halftones utilizes the small edge detection threshold  $T_{thick}$ , applying the enhancement process to the identified thick coarse edges produces low quality results. This is because detection with the small threshold  $T_{thick}$  is not very good at identifying the precise location of coarse edges. Thin edge labeling with the larger threshold  $T_{thin}$  is much more accurate in determining these locations, however, with the unfortunate tendency to miss details which is again a concern in the enhancement.

FIGS. 7A–7E are schematic diagrams illustrating an edge proximity technique for detecting missed edges during coarse edge detection on a schematic text character in accordance with one embodiment of the invention. FIG. 7A represents an exemplary text, e.g., a kanji character. When a low pass filter is applied to such a character, only the outer boundary of the letter may be detected. It should be appreciated that the inner edges may appear as high frequency structures and the low pass filtering causes the edges to be blurred and even merged with a top outer boundary. As shown in FIG. 7B, coarse edge detection with the large threshold (thin coarse edges) can miss some salient edges in region **180** and only capture outer boundary **182**, which may cause problems in edge enhancement over halftone regions. On the other hand, in FIG. 7C it can see that coarse edge detection with the small threshold (thick coarse edges) tends to mark too many pixels around edges, i.e., the edges end up being more blurred and more points are being labeled as edges with respect to the outer boundary edges, which is also a concern if enhancement is to be restricted to thick coarse edges in halftone regions. It is preferable to enhance the edges of the edge map of FIG. 7B rather than the edge map of FIG. 7C, since FIG. 7C includes points outside of the edge of the text.

In one embodiment, the edge information from the missed region **180** is determined from FIG. 7C and combined with FIG. 7B in order to accurately define all the edges. FIG. 7D illustrates the results of the application of proximity processing being applied to FIG. 7C. For proximity testing, a thick edge in FIG. 7C is examined to determine all the thick edges that have at least two thin edge strokes within a defined proximity as discussed below. This may be accomplished by identifying thick edges that are close to two or more detected thin edges or strokes (i.e., close to those passing the threshold  $T_{thin}$ ) as shown in FIG. 7D and as discussed below. The identified pixels are treated as thin coarse edges in enhancement (FIG. 7E). It should be appreciated that the results of this process only affects enhancement stages since halftone detection effectively utilizes both thin and thick coarse edges.

Around each thin coarse edge pixel, we use a temporary label to mark a symmetric ( $+/-N_{prox}$ ) rectangular proximity region of size  $(2N_{prox}+1) \times (2N_{prox}+1)$  pixels. Then, given a thick coarse edge pixel,  $N_{prox}$  pixels are checked with respect to the pixel's immediate left to determine if all pixels are marked with the temporary label. A similar check is performed on  $N_{prox}$  pixels to the immediate right, above, and below. The pixel is then activated, i.e., labeled as a "secondary" thin edge pixel, if all four proximity checks pass and the pixel itself is marked with the temporary label. The "secondary" labeling is implemented in a fashion to ensure that only the threshold-passing thin edges are used in establishing the proximity regions. Of course, in the enhancement stage secondary thin edges are treated as thin edges and undergo enhancement.

Setting  $N_{prox}$  to a large number would invoke the risk of broadening thin edges significantly on dense text regions where different text characters are at close proximity.  $N_{prox}$  is thus set to a small number,  $N_{prox}=3$  at 600 dpi, so that the procedure results in the broadening of thick edges only to the extent that helps the detection of missed salient edges (setting  $N_{prox}=3$  also provides acceptable performance at 300 dpi).

FIGS. 8A and 8B illustrate an original scanned image portion and a corresponding labeled fine scale image portion, respectively, in accordance with one embodiment of the invention. FIG. 8A corresponds to an original scanned image scanned at 600 dpi. FIG. 8B corresponds to the output of operation **138** with reference to FIG. 4A, where the white pixels represent detected edges. Fine edge detection is established through the filters of Equation (2) via:

$$D_{fh}(i,j)=E_h(i,j)*S(i,j), D_{fv}(i,j)=E_v(i,j)*S(i,j) \quad (5)$$

And there is a fine edge at pixel (i,j) in this representative channel if  $B(i,j)=1$  where

$$B(i,j)=\begin{cases} 1 & |D_{fh}(i,j)| \geq T_{fine} \text{ or } |D_{fv}(i,j)| \geq T_{fine} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

It should be appreciated that the threshold used for the examples provided herein is set to  $T_{fine}=15$  at 600 dpi ( $T_{fine}=15$  also at 300 dpi). However, any suitable threshold may be used.

Halftone detection starts by identifying all the pixels that are marked as fine edges (see FIG. 8B) but not as coarse edges (see FIG. 6B). With the symmetric, zero-delay filters utilized discussed above this determination is reduced to straightforward point-wise logic. Using Equations (4) and (6) we obtain:

$$C(i,j)=B(i,j)-(A(i,j) \text{ AND } B(i,j)). \quad (7)$$

FIG. 9 illustrates an image portion having locations marked as fine edges but not as coarse edges for the example document in accordance with one embodiment of the invention. It should be appreciated that the image represented by FIG. 9 was obtained using data shown in FIGS. 6B and 8B. While this serves as an initial estimate of the halftone dots, this process alone cannot mark halftone regions. As can be seen in FIG. 9, there are pixels on the halftone side that are not marked and there are those in the non-halftone side that are marked. The latter pixels typically correspond to locations where coarse edge detection has failed to mark detail structure on a real document edge.

FIG. 10 is a schematic diagram illustrating the counting neighborhood utilized for halftone decisions in accordance with one embodiment of the invention. In order to determine whether a pixel is over a halftone region or not, a rectangular counting neighborhood **190** is defined around the pixel and the number of locations that are marked in the initial estimate are counted. This count is compared to a threshold  $T_{cnt}$  which determines the minimum halftone density over the neighborhood. If the count is above the threshold then the pixel around which the count is established is a suspected halftone pixel, however, as will be discussed below, the constraints that affect the selection of the threshold  $T_{cnt}$  and the size and shape of the counting neighborhood need to be accounted for before a full determination.

By nature of the banded operation, the neighborhood used in the count is necessarily larger over columns than it is over



rows (this follows since the algorithm tries to store as few rows as possible). The length of this neighborhood over columns effectively determines the length of the active region with reference to FIG. 3. At 600 dpi a symmetric total length of  $N_{h,col}=2 \times 128 + 1 = 257$  pixels is utilized over the columns around the current pixel. At 300 dpi this is reduced to a total of  $N_{h,col}=2 \times 64 + 1 = 129$  pixels over the columns around the current pixel. The maximum size of the neighborhood over rows is set to 11 pixels.

Still referring to FIG. 10 the neighborhood is superimposed on various pixels in the example of FIG. 9, as well as on another example (from FIGS. 11A–11D, which shows a different portion of the same document) in order to outline three important cases that influence the counting methodology. The direction of the rows and columns is as indicated in FIG. 10. First observe that the threshold  $T_{cnt}$  cannot be set too high. As illustrated for Case 1, there are halftone pixels whose neighborhoods overlap detected coarse edges and too high a threshold will cause a miss over such pixels. Since coarse edge detection will invariably miss detail features in real edges, say over text-on-non-halftone portions of the document, counting only over pixels that are not coarse edges and using a normalization scheme runs the risk of incorrectly labeling portions of real edges over non-halftone regions as halftone. However, as seen on Case 2 and Case 3, setting a small threshold may result in a false layer of halftones around detected halftone regions. The thickness of this false layer will be particularly pronounced in column direction due to the choice of the geometry of the counting neighborhood (Case 3).

In one embodiment, the thickness of this false layer may be significantly reduced by establishing two separate counts over two partial neighborhoods around the current pixel (see FIG. 12). These two counts are also compared to a threshold  $T_{p,cnt}$  (obtained by scaling down the main threshold in proportion to the area of the partial neighborhoods) and a halftone pixel is declared if and only if the full count is over the threshold  $T_{cnt}$  and at least one of the partial counts is also over  $T_{p,cnt}$ . For the examples contained herein the column width of the partial neighborhoods is set to  $N_{p,col}=32+1=33$  at 600 dpi and  $N_{p,col}=16+1=17$  at 300 dpi. In one embodiment, all of the counting is implemented with accumulator buffers in order to avoid spurious computations.

FIGS. 11A–11D represent an original image scanned at 600 dpi, labeled coarse scale edges, labeled fine scale edges, and fine but not coarse edges, respectively. Here, FIG. 11B corresponds to a portion of FIG. 6B, FIG. 11C corresponds to a portion of FIG. 8B, FIG. 11D corresponds to a portion of FIG. 9.

Constant tones are detected by calculating the variance over two neighborhoods around each pixel. The choice of the neighborhoods is similar to the two partial counting neighborhoods used in halftone detection as illustrated with reference to FIG. 12. FIG. 12 is a schematic diagram illustrating the partial counting neighborhoods in halftone decisions in accordance with one embodiment of the invention. This allows rapid transition in detection from other types of regions to constant-tone regions and vice versa. A constant-tone pixel may be declared if the variance over either of the neighborhoods is below a threshold  $T_{const}$ . The column width of the partial neighborhoods is set to  $N_{c,col}=64+1=65$  at 600 dpi and scaled suitably for 300 dpi. The row width is the same as the halftone detection stage and it is set to  $N_{c,row}=11$ . The threshold  $T_{const}$  may be loosely dependent on scanner characteristics and it is set to  $T_{const}=9 \times 9 = 81$  at both resolutions. It should be appreciated

that there may be many device specific factors involved in bleed-through removal processing, e.g., color of the scanner backplane, thickness of paper, etc.

Of course, establishing variance calculations over rectangular regions does not allow constant tones to be detected inside densely populated text regions. Therefore the above detection is followed with an extension technique where non-halftone and non-edge pixels that are connected to an identified constant tone pixel and that are close in color (within  $(T_{const})^{1/2}$ ) are also marked as constant-tones. This enables seamless extension of constant-tones into text regions. However, as will be further discussed below, while regions around text characters will be correctly identified, the holes inside characters may be missed. The effects of this problem are negligible as discussed below.

For each pixel the assigned independent labels are first used to identify a single label for each color channel using Table 2 (in Table 2 the “\*” indicates a don’t care condition).

TABLE 2

HALFTONE	COARSE EDGE (THIN)	FINE EDGE	CONST. TONE	LABEL
1	1	*	*	COARSE EDGE
1	0	*	*	HALFTONE
0	1	0	*	CONDITION 1
0	1	1	*	CONDITION 2
0	0	1	*	FINE EDGE
0	0	0	1	CONSTANT TONE
0	0	0	0	OTHER

These individual labels are then combined to arrive at a final label. It should be appreciated that the thin coarse edges are used in determining the coarse edge labels for each color channel and therefore the final overall label as discussed above. The CONDITION 1 and 2 labels in Table 2 are used as discussed below in order to accommodate cases where a halftone is detected in one channel and a coarse edge in another (such as the case of green text over a red halftone region).

The final label that combines the labeling information for each channel is given as COARSE EDGE label if any channel is marked COARSE EDGE or if one channel is marked HALFTONE and at least one other marked CONDITION 1 or 2. Otherwise, the final label is HALFTONE if any channel is marked as HALFTONE, else FINE EDGE if any channel is marked FINE EDGE or CONDITION 2, else CONSTANT TONE if all channels are marked CONSTANT TONE. If none of these conditions hold then the final label is OTHER. For the discussion below it is assumed that the final label for pixel (i,j) is given as X(i,j).

With respect to the embodiment where multiple filtering branches are used, e.g., FIG. 4C, two additional labels, referred to as CONDITION 3 (CONDN3) and CONDITION 4 (CONDN4), are used for the identification of low resolution halftone regions. The incorporation of the CONDITION 3 and CONDITION 4 labels are algorithmically described below. It should be appreciated that LOW\_RES\_HT and LOW\_RES\_COARSE\_EDGE are marked as described above in the earlier algorithm in each channel, independently of other labels. The below description provides an overview of the incorporation of CONDN3 and CONDN4 labels:



---

```

If at least one channel is marked as LOW_RES_HT, then we set CONDN3 true.
  If at least one channel is marked as LOW_RES_COARSE_EDGE, then we set
  CONDN4 true.
  Then after the combined labels are assigned by the earlier algorithm as discussed above
  with reference to TABLE 2, we use CONDN3 and CONDN4 to modify these labels as follows:
  if CONDN3 is true then
    if CONDN4 is false then
      if the earlier final label is not HALFTONE then
        if the earlier final label is not CONSTANT TONE
          then
            mark all channels as LOW_RES_HALFTONE
            and do resolution dependent
            enhancement/blurring on them.
          endif
        endif
      endif
    endif
  endif

```

---

(It should be appreciated that all implied else branches above do not change the labels assigned by the earlier algorithm.)

FIGS. 13A and 13B are schematic diagrams illustrating the per-pixel labels obtained for the examples depicted in FIG. 5A and FIG. 11A, respectively. Here, black regions 192 correspond to fine edges, dark gray regions 193 correspond to coarse edges on halftones, gray regions 194 correspond to halftones, light gray region 195 correspond to constant tones, and white regions 196 correspond to pixels labeled other. It should be appreciated that the main segmentation of the various regions is accomplished correctly with some small mistakes due to various limitations discussed above. As can be seen, the halftone segmentation in FIG. 13A is more accurate than in FIG. 13B, which contains a false layer of halftone at the top of the halftone region. This is due to the structure of the halftone detection neighborhoods discussed above, and as can be seen, it is more pronounced in the column direction than it is in row direction. With the partial neighborhoods set to  $N_{p,col}=33$  pixels in 600 dpi as discussed above, the adverse effects of the false layer is negligible in conjunction with the error tolerant enhancement procedures discussed with reference to the enhancement and descreening stages. As mentioned above, the constant-tone regions do not extend to holes inside characters with such locations erroneously labeled as OTHER (for example see the inside of the text character "p" in FIG. 13A). Further examples, and properties are described in more detail below.

The enhancement and descreening stage primarily takes into account that the output device is a printer and targets an observation model that is more sensitive to low-pass frequency characteristics especially for the enhancement of coarse edges. However, more elaborate techniques may be used for other scenarios. Enhancement proceeds on color channels independently except for final normalization, where the computed enhancement differentials are jointly scaled to ensure that pixel values in each color channel remain within maximum and minimum limits.

Some labeling mistakes are inevitable, accordingly the embodiments described herein combine simple enhancement with the labeling to achieve very robust performance. For example, the addition of an extra branch in FIG. 4C provides the capability to detect low resolution halftones in order to apply weak enhancement to the detected low resolution halftones as described below with reference to FIG. 24. A property that enables this robustness is the very similar methods used in the enhancement of coarse and fine edges. Another property that is emphasized in the design is

the sharing of computations between the labeling stage described with reference to FIG. 4A and the enhancement stage. For example, the low pass filtered data used in the labeling stage is reused in descreening and enhancement in conjunction with the detected labels. These factors contribute to a robust overall design of very low computational and memory complexity. For the discussion below,  $P(i,j)$  denotes the representative channel in the output document.

Descreening is accomplished by inputting the already low pass filtered input values at all the pixels that have been labeled as halftone (See FIG. 14), i.e., using equation (1) yields:

$$P(i,j)=L(i,j), \text{ if } X(i,j)=\text{HALFTONE.} \quad (8)$$

It should be appreciated that this incurs minimal complexity.

At fine edges an unsharp masking procedure, which is generally known, is used by computing a sharpening term that is given by the input pixel value minus the already computed low pass filtered pixel value, i.e., using equation (1) yields:

$$P(i,j)=S(i,j)+\lambda(S(i,j)-L(i,j)), \text{ if } X(i,j)=\text{FINE EDGE.} \quad (9)$$

This again incurs minimal complexity. The sharpening term is multiplied with a sharpening factor and then added to the original pixel value to determine the enhanced output at pixels labeled as fine edges (See FIGS. 14A and 14B which illustrate the original scanned document and the enhanced output, respectively).

The enhancement of coarse edges is very similar to the enhancement of fine edges discussed above and is also carried out using unsharp masking. However, a small (3x3) low pass filter  $K(i,j)$  is utilized on the input to obtain:

$$l(i,j)=K(i,j)*S(i,j)$$

The sharpening term is obtained as this value minus the already computed low pass filtered pixel value. The sharpening term is again multiplied with a sharpening factor and then added to the low pass filtered value to determine the enhanced output at pixels labeled as coarse edges, thereby yielding:

$$P(i,j)=L(i,j)+\lambda(l(i,j)-L(i,j)), \text{ if } X(i,j)=\text{COARSE EDGE.} \quad (10)$$

where we have again used equation (1). For 600 dpi results in this paper the small low pass filter kernel is set to



$$K(i, j) = \begin{bmatrix} .25 \\ .50 \\ .25 \end{bmatrix} \times [.25 \ .50 \ .25] \quad (11)$$

It should be appreciated that this filter should be tuned for optimum performance on the output printer.

As can be seen on the enhancement results in the halftone region with reference to FIGS. 14A and 14B, the described enhancement process on coarse edges results in some visibility of halftone dots close to the edges. In one embodiment, this is intentionally done to generate sharp output since halftone screened printouts of these enhanced results remove this visibility and further processing, e.g., the thinning of edges, and more sophisticated enhancement is not necessary.

The bleed-through removal method is then applied and is meant as an exemplary use of the constant-tone labeling. The bleed-through removal method operates on constant-tone regions and starts by calculating a histogram of all the constant-tone pixels in the processed row (this is essentially at the top row stored by the algorithm). In order to avoid higher dimensional histograms, the histogram is computed on a luminance-like channel or index obtained by averaging the pixel values in each color channel. For each bin in the histogram, the average color of all pixels falling into that bin is stored. In one embodiment, simple checks are performed to ensure widely varying colors do not fall into the same bin due to luminance-like indexing. In order to ensure continuity this histogram is further averaged with a stored "previous" histogram (and its average colors), and the results are utilized as well as stored for the processing of the next row.

Histogram segmentation is then performed and each segment in the histogram is assigned a single representative color given by the conditional average color. The colors of the constant-tone pixels in the enhanced image are tested against these representative colors and those pixels with colors that are within a certain threshold (set to  $2((T_{const})^{1/2})$ ) of the corresponding segment's color are assigned the representative color.

Of course, this technique is mainly effective in removing bleed-through artifacts over regions marked as constant-tone. However, it should be appreciated that at the expense of increased memory, constant-tone labeling on halftone regions may be performed as well and bleed-through removal may be carried out on de-screened data. With this simple technique, light bleed-through in extended constant-tone regions is substantially eliminated in the input, where artifacts are the most visible. By tuning thresholds it is possible to combat more aggressive artifacts provided that the histogram segmentation can be done accurately.

The simulation examples described below provide exemplary applications of the labeling approach embodiments described herein. FIGS. 15A and 15B are exemplary schematic diagrams illustrating the robustness of the algorithm in detecting different types of halftones. More specifically, FIGS. 15A and 15B include a series of images illustrating the application of the embodiments described above for detecting different types of halftones in accordance with one embodiment of the invention. The original scanned image for both FIGS. 15A and 15B, 210-1 through 210-4 is a document scanned at 600 dpi containing four halftoned sub-images. As illustrated in box 210 of FIG. 15A, each of four sub-images 210-1a through 210-4a uses a different

halftone screen. Image portions 210-1 through 210-4 are scaled to 25% of the original image. In FIG. 15B it should be appreciated that that the labeling obtained by the algorithm is virtually insensitive to the changing halftone screens, and each of the four sub-images 210-1 through 210-4 is labeled in a similar fashion. The main results of the labeling are correct except for the regions identified as Case 1 216 and Case 2 218. Case 1 216 is the situation of a false layer of halftones due to the structure of the halftone detection neighborhoods which has been discussed above. It should be appreciated that Case 1 situations have the potential to give rise to adverse effects if the layer is overlapping fine edges. However, as will be shown, the utilized enhancement stage is robust and minimizes such effects. Adverse effects associated with Case 1 situations are similar to Case 3 situations that are discussed in detail below. Case 2 situations, as identified in FIG. 15B and also in FIG. 16B correspond to the situation of very dark or very light regions inside halftone regions. Such regions do not contain halftone dots and they will be missed if they are of appreciable size since the algorithm strives to keep the relevant neighborhood sizes to a minimum for the above stated complexity, memory, and memory-bandwidth reasons. Of course, Case 2 situations have virtually no effect on the enhanced output since no halftone dots are present to cause artifacts. However, it will be apparent to one skilled in the art that Case 2 situations may be corrected by region analysis algorithms as these errors typically manifest themselves as constant-tone regions surrounded by detected halftones.

FIGS. 16A through 16E, 17A through 17E, and 18A through 18E are exemplary schematic diagrams illustrating the performance of the above described embodiments of the invention on full pages of input. Each one of the examples are chosen to illustrate how the technique handles a challenging compound document type. In each Figure set, the corresponding Figure designations signify the following: (A) show the original, (B) show the labeling obtained by the algorithm, (C) show hand segmented (ground truth) halftone and non-halftone segmentation, (D) show segmentation obtained by the algorithm, and (E) show the enhanced output. The documents are intended to illustrate the main behavior of the algorithm on full pages of input and they are scaled to 25% for display purposes. Example halftone dots from each document are shown on the right bottom of the (A) set of Figures. Each of the original documents are scanned at 600 dpi.

The halftone portions in each document can be seen in detail on the (C) set of Figures which are the hand-segmented ground truth separation of the page into halftone and non-halftone. In order to calculate quantitative detection probabilities with respect to the ground truth, the halftone and non-halftone segmentation of the documents are shown as obtained by the algorithm in the (D) set of Figures. Of the five labels outputted by the algorithm "COARSE EDGE" and "HALFTONE" constitute pixels over halftone regions whereas the remaining three labels in the X field denote non-halftone pixels. By comparing the (C) and (D) sets of corresponding Figures it should be appreciated that the algorithm detects varying halftones with success.

FIGS. 16A-16E illustrate performance of the embodiments described above on a multilingual document containing halftones and fine text. For this document it is important to identify halftone regions as accurately as possible since mistakes may lead to loss of quality on fine text. Observe that the document includes fine text with detail strokes and the document is labeled correctly except for the identified problems as discussed below with reference to FIG. 16B. In



addition to earlier identified Case 1 and Case 2 situations, this document allows for discussion of Case 3 situations where some fine text close to halftone regions is labeled as being inside the halftone regions (see corresponding regions in FIGS. 16(B), (C), and (D)). The enhancement ramifications of this situation are discussed in more detail with reference to FIG. 19. Case 3 situations typically arise when text characters containing very detailed strokes are close to halftone regions. The detailed strokes are missed by the coarse edge detector but not by the fine edge detector, and they thus enter the halftone counting process. The overall effect is that the main halftone counting neighborhood overlaps legitimate halftone regions and the main count passes  $T_{cnt}$ . At the same time the partial neighborhoods count the missed fine detail, sometimes passing  $T_{p,cnt}$  resulting in false halftone decisions. Note however that this is only a problem if fine text is present very close to halftone regions. The final issue with respect to FIGS. 16A–16E is a combination of Case 2 and Case 4 as illustrated in FIG. 16B. Discussion of Case 4 situations is more clearly illustrated with reference to FIGS. 17A–17E.

FIGS. 17A–17E illustrate performance of the embodiments described above on a multilingual document having a very irregularly shaped halftone region. Observe in FIGS. 17(B), (C), and (D) that most of the halftone is correctly detected, and in particular, our algorithm which only operates over small neighborhoods is very successful in detecting the global segmentation. The main problem with FIG. 17 is illustrated in Case 4 (FIG. 17B) and corresponds to the situation where very small halftone regions are missed by the algorithm. Since the neighborhood sizes used by the algorithm are very small, the size of these missed regions are also expected to be small. As such, they typically pose insignificant problems in enhancement.

FIGS. 18A–18E again shows performance of the embodiments described above on a document with an irregularly shaped halftone region. The document of FIGS. 18A–18E has a smooth transition or gradation from non-halftone to halftone and vice versa (above and below the middle halftone region). The technique described herein not only detects the halftone regions correctly irrespective of the gradation, but the enhancement results on this document are also very robust to this effect. FIG. 18 figure contains Case 2 situations in the dark regions (not marked) but the main potential problem is as shown in Case 5 where there is fine text in a halftone to non-halftone transition area. Many techniques are rigid in their segmentation and enhancement decisions and such areas typically result in very visible abrupt transitions in their output. However, as discussed below with respect to enhancements this situation is not the case for the techniques discussed herein.

In order to discuss quantitative behavior let  $P_{D,h}$  denote the probability that the algorithm detects a halftone pixel given that the ground truth labeling for the pixel is halftone. Similarly let  $P_{D,n}$  denote the probability that the algorithm detects a non-halftone pixel given that the ground truth labeling for the pixel is non-halftone. The conditional false alarm or miss probabilities are one minus these probabilities. The overall detection probability  $P_D$  is then given by  $P_{D,h}$  times the number of ground truth halftone pixels plus  $P_{D,n}$  times the number of ground truth non-halftone pixels. For the results in FIG. 16D  $P_{D,h}=0.985$ ,  $P_{D,n}=0.960$ , and  $P_D=0.970$ . In FIG. 17D,  $P_{D,h}=0.974$ ,  $P_{D,n}=0.986$ , and  $P_D=0.984$ . Finally in FIG. 18D  $P_{D,h}=0.977$ ,  $P_{D,n}=0.978$ , and  $P_D=0.978$ . It is clear that these are very good results for the simple technique embodied herein.

Finally note that the labeling stage is designed to be used in conjunction with the enhancement in a very low complexity environment. Many of the identified issues, as well as the mistakes over very small regions, can be corrected if one is willing to resort to region completion and other more elaborate techniques known to one skilled in the art.

The (E) figure set in FIGS. 16, 17, and 18 illustrate the enhancement performance on the earlier examples. These Figures are illustrative that there are no catastrophic errors. Note that the enhancement results include no color or gamma correction to accommodate the scanner characteristics. In one embodiment, the scanner and printer color corrections are handled jointly within the printing pipeline in order to gain computational savings as discussed above. Thus, the presented results are not meant to be judged in terms of the correctness and appearance of colors. The (E) set figures also do not include bleed-through removal which will be discussed below.

FIGS. 19A–C, 20A–C, and 21A–C show detailed enhancement performance on the earlier identified potential problem areas. FIGS. 19A–C show the enhancement results under situations due to Case 1 and Case 3 type errors (e.g., compare FIG. 19 to the lower right of the halftone region that is at top left in FIG. 16). These errors will only affect enhancement in situations where there is fine text close to halftone regions. As can be seen in FIGS. 19–21 even in such cases, thanks to the method utilized in enhancing coarse edges, significant problems are avoided. The similarity in the way enhancement factors are calculated discussed above prevents significant visible differences. However, there may be rare occurrences where portions of fine text strokes in such regions may appear blurred if such strokes are not detected in coarse edge detection.

FIGS. 20A–20C illustrate results over Case 4 situations. In this situation small halftone regions, depending on their alignment with respect to the counting neighborhoods in halftone detection, may be missed. The visual effects of such mistakes are negligible since the missed regions are very small to begin with (compare FIG. 20 to the same region in the lower left of FIG. 17 to see the relative size on the full page document).

FIGS. 21A–21C show the performance of the technique over Case 5 problems as seen with reference to FIG. 18B. It should be appreciated that the enhancement in FIG. 21B is uniform and the transition in halftones is correctly handled. Again, the similarity in the way coarse and fine enhancement is designed allows us to avoid noticeable problems. Any visible issues are even further subdued on actual printouts.

FIGS. 22A–22C are a schematic diagrams illustrating the technique's performance on bleed-through removal. It should be noted that the bleed-through is successfully removed and the missed constant tones, pixels labeled as "OTHER" inside and around text regions, do not cause visible problems in the enhanced output. In FIGS. 22A–C the left hand column depicts the originals, the middle column depicts the enhanced documents with bleed-through removal and the right column shows the obtained labels. It should be appreciated that FIG. 22A corresponds to a portion of FIG. 16, FIG. 22B corresponds to a portion of FIG. 17, and FIG. 22C corresponds to a portion of FIG. 18.

FIG. 23A is a simplified schematic diagram illustrating a portion of a compound document which may be mistakenly characterized as a halftone region. With just one low pass filter branch, as illustrated in FIG. 4A, lines 246 of FIG. 23A are detected as halftone regions in section 242 and not enhanced, while the portion of lines 246 within region 240



are detected as non-half-tone areas and consequently strongly enhanced. Accordingly, a large artifact may be observed between section 242 and section 240.

FIG. 23B is a simplified schematic diagram illustrating a scheme to address the mistaken characterization discussed with reference to FIG. 23A, through edge detection at three scales in accordance with one embodiment of the invention. An exemplary schematic illustrating edge detection at three scales is illustrated with reference to FIG. 4C. Performing edge detection using at least three scales enables a third region to be introduced in between strong enhancement region 248 and blurred region 252. The third region, weak enhancement region 250, smoothes the transition from blurred region 252 to strong enhancement region 248. Therefore, the artifact is not as observable. With reference to FIG. 4C, low pass filter A 134', provides increased blurring relative to low pass filter B 134 in one embodiment. Thus, low pass filter A 134' is adapted for identification of low resolution half-tone regions, while low pass filter 134 is better suited for high resolution half-tone regions. It should be appreciated that while FIG. 4C illustrates two additional branches 133 and 133' in addition to unfiltered branch 135, the scheme described herein may include any suitable number of filtering branches in excess of two.

It should be further appreciated that the relative locations of lines 254 and 256 may be controlled through the bandwidth of the corresponding low pass filters. That is, low pass filter A 134' of FIG. 4C impacts the position of line 254 of FIG. 23B. Similarly, low pass filter B 134 of FIG. 4C is associated with the position of line 256 of FIG. 23B. In one embodiment, as the bandwidth of low pass filters 134 and 134' increase, the corresponding locations of line 256 and 254, respectively, move toward the right (from strong enhancement region 248 to blurred region 252). One skilled in the art will appreciate that as additional branches are added to FIG. 4C, additional regions will be included in FIG. 23B. That is, while regions 248, 250, and 252 are shown in FIG. 23B, the number of regions inserted between enhancement region 248 and blurred region 252 may be further increased by the addition of extra branches to FIG. 4C.

FIG. 24 is an alternative representation of the effect of an additional branch being included in the labeling module as described with reference to FIG. 4C, in accordance with one embodiment of the invention. Here, images 260a through 260n, are sequentially placed in order of relative half-tone frequency. That is, image 260a is associated with the highest half-tone frequency and the half-tone frequencies successively decrease from image 260a to image 260n. In other words, the distance between half-tones successively decreases as illustrated by distances  $d_1$ – $d_5$ . Here, strong blurring provided by low pass filter A ( $LPF_A$ ), is applied to images 260a through 260e. The low-resolution half-tones associated with images 260d and 260e may not be detected as a half-tone region by low pass filter B ( $LPF_B$ ), thereby eliminating the ability to apply weak enhancement. That is, by adding low pass filter A, the low resolution half-tones of images 260d and 260e are labeled as such and weakly enhanced, i.e., enhanced or blurred according to an associated resolution, in order to buffer the transition between strong blurring and strong enhancement depicted along line 254 moving in the direction from image 260a to image 260n.

Still referring to FIG. 24, images 260a–260c are associated with serious Moiré artifacts if the image data is not filtered prior to enhancement. Downstream from image 260e there are no Moiré artifacts. With this technique it is possible to suppress serious Moiré artifacts, do aggressive enhancement downstream from 260e, and maintain a transition

region of resolution dependent enhancement/blurring that avoids transition artifacts and maximizes visual quality. It should be appreciated that the images associated with FIGS. 15A and 15B illustrate the successive reduction in half-tone resolution depicted by images 260a–260d of FIG. 24. Thus, image 210-1 of FIG. 15A may be considered a low resolution half-tone image, similar to images 260d and 260e of FIG. 24. Thus, filter B 134 of FIG. 4C may not detect the low resolution half-tones, however, filter A 134' is configured to enable detection of the low resolution half-tones so that weak enhancement, which may be referred to as resolution dependent enhancement/blurring, may be applied to the low resolution half-tones as discussed above. The resolution dependent enhancement component can be accomplished by using different sharpening factors in the enhancement scheme described above. The resolution dependent blurring component can be accomplished by utilizing different low pass filters in the Moiré removal scheme described above.

The significant computational complexity required by the algorithm is summarized below for each pixel in each color channel. All operations are on integers with most operations on 8 or fewer bit integers.

The symmetric, separable low pass filtering described above with reference to FIGS. 4A and 4C, needs an add and a multiply for about half the number of taps over rows and over columns (due to the separable filter) for each pixel in each channel in one embodiment. For the mentioned parameters at 600 dpi, this corresponds to  $2 \times (3+1) = 8$  additions and multiplies for each pixel in each channel (a final normalizing division or truncation is needed in integer operation). Of course the multiplies are for full resolution operation and the filter can be implemented multiplier-less as well.

Edge detection requires filtering with the filters of Equation (2) and thresholding. The filters are separable and much simpler than the low pass filter. Filtering requires 1 shift, 2 additions, and 1 subtraction. Two filtering computations are carried out per pixel for each color channel and another factor of two comes in for coarse in addition to fine edge detection, yielding 4 shifts, 8 additions, and 4 subtractions.

For the computations discussed with reference to thin coarse edges, an average of  $2N_{prox}+1$  pixels need to be marked with the temporary label in the worst case for each pixel in each channel. Proximity checks using counts over rows and columns are implemented here. A single row buffer is utilized that stores the count of temporary labels on each column corresponding to the sum over  $2N_{prox}+1$  rows. In this fashion the horizontal or vertical proximity check may be reduced to 2 checks for labels and corresponding increment by 1 or decrement by 1 operations (a possible increment to update the buffer for a new row and a possible decrement to subtract the influence of an out of scope row). Hence, in the worst case, the total becomes 4 checks for labels and 4 increment/decrement operations for each pixel in each channel.

In one embodiment, the counting required for half-tone detection is implemented in two stages as counts over  $N_{h,row}$  rows in each column followed by a count over columns. Similar to proximity check calculations, a one row buffer is utilized to store the counts over rows for each column resulting in two increment/decrement operations and two label checks. Counting over columns add two addition/subtraction operations which is carried out for the three counting neighborhoods as discussed with reference to half-tone detection and FIG. 10. Therefore, a total of two increment/decrement and  $3 \times 2 = 6$  addition/subtraction operations for each pixel in each channel is incurred.



Edge enhancement requires 2 additions and 1 multiplication, which can be done with via a shift depending on  $\lambda$ . In the enhancement of coarse edges, the filter of Equation (11) requires 2 additions and a shift over rows as well as columns. Hence, worst case edge enhancement needs 6 additions, 2 shifts, and 2 multiplies per edge pixel, in each color channel. Depending on the normalization method a normalizing division may also be performed for each edge pixel, in each color channel.

It can be seen from the above that the main portion of the algorithm requires on the order of 10 multiply and 30 addition/subtraction operations for each pixel in each channel. The bulk of this computation comes from low pass filtering and edge detection. It should be appreciated that any technique that accomplishes descreening on input data needs to be prepared for per pixel low pass filtering. Similarly, for even modest quality in a copier pipeline, at least some edge detection and enhancement is necessary. Hence, the computational resources added by our algorithm over a "baseline" algorithm are those needed by a secondary edge detection and halftone detection. This is on the order of 10 addition/subtraction and 6 increment/decrement operations and corresponds to about 25% increase in computation over the most naive scanner-side algorithm in a copier pipeline.

For applications requiring bleed-through removal, constant-tone detection is also implemented with a row buffer that stores calculations over columns. Squares for variances, etc., are implemented via lookup tables resulting in 1 add and 1 subtract each for the mean and the variance. Calculations for each neighborhood add 1 add and 1 subtract each for the mean and variance followed by a mean normalizing multiply and a subtraction. The total computation comes to  $3 \times (2 \times 2) = 12$  addition/subtraction operations plus 2 multiplies and 2 subtractions for each pixel in each channel. The luminance-like index calculation in the simple bleed-through removal scheme described above requires 2 additions and a normalization per pixel. Histogram accumulations further require an addition per pixel, for each color channel. Final histogram normalization, averaging, and segmentation are carried out on a row by row basis. Thus, carrying out the bleed-through scheme further requires on the order of 20 addition/subtraction and 3 multiplication operations. In summation, the technique described herein, requires about 60% more computation over the most naive scanner-side algorithm in a copier pipeline, in one embodiment.

It should be further appreciated that additional computations in the copier pipeline, e.g., color correction and halftoning require significant computation. Hence, the practical computational complexity added by the incorporation of the above described technique is typically an even smaller percentage of the overall computational complexity. One skilled in the art will appreciate that the use of three extra row buffers mentioned above, which are included for computational purposes, are not required in all implementations. However, considering these buffers raises the memory requirements of the technique by another 0.8%, thereby bringing the total memory requirements to approximately 1.5% of full page memory.

In summary, the embodiments described herein provide a multi-resolutional, single pass technique that accomplishes segmentation and enhancement of compound documents under relatively low computational complexity and memory requirements. The technique is scalable and provides high halftone detection probabilities. The enhancement stage is designed to operate jointly with the labeling stage, and provide robust operation even under situations where label-

ing errors may occur. The banded architecture and the sharing of computations between stages enable the technique to provide significant functionality at low computational, memory, and memory bandwidth complexity.

The multi-resolutional technique does processing on the input (fine scale) and the low pass filtered version of the input (coarse scale). Edge detection is done at the two resolutions to mark "coarse edges" on the low pass filtered image and "fine edges" on the original image (See FIGS. 4A and 4C for a simplified schematic of the labeling process). With the low pass filter chosen to suppress most of the halftone, edge detection on the low pass filtered image marks most of the real edges (due to image edges, text, graphics, etc.), while fine edge detection marks the real edges and halftone dots. All locations that are marked as fine edges but not as coarse edges serve as initial estimates of halftone dots (i.e., result C in FIGS. 4A and 4C). Such locations are counted around each pixel and the result thresholded in order to determine halftone pixels. Coarse edges on detected halftone regions and fine edges on non-halftone regions make up two of the labels in the X field. The variance around each non-halftone pixel is calculated to determine if that pixel is over a constant-tone region. In one embodiment, five possible labels as identified above in Table 1 are possible. Once labeling is done, edge enhancement is carried out on detected coarse and fine edge pixels and descreening is accomplished by low pass filtering the halftone pixels. If desired, simple bleed-through removal is applied. It should be appreciated that computations in the enhancement stage share some of the computational results of the labeling stage in order to reduce the computational complexity.

Consequently, the embodiments described herein combines segmentation and region classification into one seamless labeling stage that is accomplished jointly with descreening and enhancement in a single pass algorithm of comparatively negligible computational complexity, memory, and memory-bandwidth requirements. The technique is fully automated, it is applicable to general compound documents, and it may handle a multitude of halftone types over regions of arbitrary sizes and shapes. While the term "halftone dots" is used throughout it should be appreciated that the technique may be used to detect other types of halftones as well. The above described embodiments may be used for multilingual documents and are effective on compound documents with text characters containing fine strokes.

With the above embodiments in mind, it should be understood that the invention may employ various computer-implemented operations involving data stored in computer systems. These operations are those requiring physical manipulation of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. Further, the manipulations performed are often referred to in terms, such as producing, identifying, determining, or comparing.

The invention can also be embodied as computer readable code on a computer readable medium. The computer readable medium is any data storage device that can store data which can be thereafter read by a computer system. The computer readable medium also includes an electromagnetic carrier wave in which the computer code is embodied. Examples of the computer readable medium include hard drives, network attached storage (NAS), read-only memory, random-access memory, CD-ROMs, CD-Rs, CD-RWs, magnetic tapes, and other optical and non-optical data



storage devices. The computer readable medium can also be distributed over a network coupled computer system so that the computer readable code is stored and executed in a distributed fashion.

Any of the operations described herein that form part of the invention are useful machine operations. The invention also relates to a device or an apparatus for performing these operations. The apparatus may be specially constructed for the required purposes, or it may be a general purpose computer selectively activated or configured by a computer program stored in the computer. In particular, various general purpose machines may be used with computer programs written in accordance with the teachings herein, or it may be more convenient to construct a more specialized apparatus to perform the required operations.

The above described invention may be practiced with other computer system configurations including hand-held devices, microprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, main-frame computers and the like. Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

What is claimed is:

**1.** A method of labeling digital image data associated with a compound document in a single pass, comprising:

receiving image data;  
 distributing the image data to multiple branches, wherein at least two of the multiple branches include filters configured to blur the image data at different resolutions;  
 detecting edges of the blurred image data at the different resolutions;  
 detecting edges of unfiltered image data;  
 combining data output from the at least two of the multiple branches; and  
 determining whether the combined data is to be labeled as a low resolution halftone.

**2.** The method of claim **1**, wherein the method operation of detecting edges of the blurred image data at the different resolutions includes,

suppressing halftones associated with the image data at a first suppression level in a first branch of the at least two of the multiple branches; and  
 suppressing the halftones at a second suppression level in a second branch of the at least two of the multiple branches.

**3.** The method of claim **1**, wherein the method operation of combining data output from the at least two of the multiple branches includes,

performing a logical AND operation with output from the at least two of the multiple branches to identify an edge.

**4.** The method of claim **1**, wherein a number of the multiple branches is equal to three.

**5.** The method of claim **1**, wherein, the method operation of determining whether the combined data is to be labeled as low resolution halftone data includes,

counting a number of logical values associated with real edges;  
 comparing the number with a threshold value;  
 if the number is greater than or equal to the threshold value, then the method includes,

labeling a region associated with the combined data as a low resolution halftone region.

**6.** A method of segmenting a compound document for enhancement during replication of the compound document, comprising:

filtering data representing a portion of the compound document through a filter associated with a first resolution;  
 filtering the data representing the portion of the compound document through a filter associated with a second resolution;  
 detecting edges on both, an output of the filter associated with the first resolution and an output of the filter associated with the second resolution;  
 combining data representing detected edges from both outputs;  
 subtracting the combined data from the output of the filter associated with the second resolution; and  
 determining whether a pixel corresponding to the data representing the detected edges is over a halftone region based upon a result from the subtracting.

**7.** The method of claim **6**, wherein the method operation of determining whether pixels corresponding to the data representing the detected edges are over a halftone region includes,

counting a neighborhood of pixel values around the pixel;  
 and  
 thresholding a result of the counting.

**8.** The method of claim **7**, wherein the method operation of counting a neighborhood of pixel values around the pixel includes,

defining two subsets of the neighborhood; and  
 separately counting pixel values associated with each of the two subsets.

**9.** The method of claim **8**, further comprising:  
 scaling down a main threshold proportionately to an area of a corresponding subset of the neighborhood; and  
 comparing results of the separate counting to corresponding scaled thresholds.

**10.** The method claim **6**, further comprising:  
 detecting edges on unfiltered data;  
 combining data representing edges from the unfiltered data and the output of the filter associated with the second resolution; and  
 identifying halftone regions from the combination of the data representing edges from the unfiltered data and the output of the filter associated with the second resolution.

**11.** The method of claim **6**, wherein the first resolution causes a more pronounced blurring effect relative to the second resolution.

**12.** A computer readable medium having program instructions for segmenting a compound document for enhancement during replication of the compound document, comprising:

program instructions for filtering data representing a portion of the compound document through a filter associated with a first resolution;  
 program instructions for filtering the data representing the portion of the compound document through a filter associated with a second resolution;  
 program instructions for detecting edges on both, an output of the filter associated with the first resolution and an output of the filter associated with the second resolution;  
 program instructions for combining data representing detected edges from both outputs;



31

program instructions for subtracting the combined data from the output of the filter associated with the second resolution; and

program instructions for determining whether a pixel corresponding to the data representing the detected edges is over a halftone region based upon a result from the subtracting.

13. The computer readable medium of claim 12, wherein the program instructions for determining whether pixels corresponding to the data representing the detected edges are over a halftone region includes,

program instructions for counting a neighborhood of pixel values around the pixel; and  
program instructions for thresholding a result of the counting.

14. The computer readable medium of claim 13, wherein the program instructions for counting a neighborhood of pixel values around the pixel includes,

program instructions for defining two subsets of the neighborhood; and  
program instructions for separately counting pixel values associated with each of the two subsets.

15. The computer readable medium of claim 12, further comprising:

program instructions for detecting edges on unfiltered data;

program instructions for combining data representing edges from the unfiltered data and the output of the filter associated with the second resolution; and

program instructions for identifying halftone regions from the combination of the data representing edges from the unfiltered data and the output of the filter associated with the second resolution.

16. The computer readable medium of claim 12, wherein the first resolution is configured to cause greater blurring of unfiltered data relative to the second resolution.

17. An image replication system, comprising:

a labeling module configured to segment image data corresponding to a compound document, the labeling module including,

multiple branches including edge detection modules for detecting edges of the compound document, a plurality of the multiple branches capable of filtering the image data at different resolutions;

logic for combining output of at least two of the multiple branches and subtracting the combined output from output of one of the at least two multiple branches in order to initially identify a portion of the compound document as a halftone portion; and

a counting module configured to count pixel values within a neighborhood to determine whether the initially identified halftone portion is finally labeled as a halftone region.

18. The image replication system of claim 17, wherein the counting module includes,

a threshold module configured to compare a pixel counting result to a threshold value, wherein if the threshold value is less than or equal to the pixel counting result, then the initially identified halftone portion is finally labeled as the halftone region.

19. The image replication system of claim 17, wherein the multiple branches include a single branch capable of performing edge detection on unfiltered data.

20. The image replication system of claim 17, wherein the labeling module further includes:

32

a labeling module configured to further label the halftone region as one of a text on halftone region and a no-text on halftone region.

21. The image replication system of claim 17, further comprising:

a descreening module configured to descreen data associated with the halftone region that is received from the labeling module, the descreening module further configured to output the filtered image data corresponding to one of the different resolutions.

22. The image replication system of claim 21, further comprising:

an enhancement module configured to enhance edges through an unsharp masking scheme.

23. The image replication system of claim 22, wherein the halftone portion is a low resolution halftone portion, the low resolution halftone portion being defined as a halftone region without edges when processed through a branch without filtering capability and a halftone region having edges when processed through one of the plurality of the multiple branches capable of filtering.

24. The image replication device of claim 17, wherein the image replication device is a device selected from the group consisting of a copier, a printer, a scanner, and a facsimile.

25. An integrated circuit capable of segmenting image data corresponding to a compound document, comprising:

labeling circuitry configured to identify an image data region type, the labeling circuitry having multiple branches, each of the multiple branches including,

filter circuitry capable of filtering the image data, each of the multiple branches associated with a different filter resolution; and

edge detection circuitry configured to detect edges of the filtered image data;

the labeling circuitry further including, circuitry for combining at least two outputs of the multiple branches and subtracting the combined output from output of one of the at least two outputs of the multiple branches; and  
circuitry for identifying the image data region type based upon a value defined through the circuitry for combining the at least two outputs of the multiple branches.

26. The integrated circuit of claim 25, wherein the filter circuitry is configured to low pass filter the image data in order to blur halftones prior to edge detection.

27. The integrated circuit of claim 25, wherein the circuitry for combining at least two outputs of the multiple branches includes,

circuitry for adding the at least two outputs; and  
circuitry for subtracting a result of the circuitry for adding from one of the at least two outputs.

28. The integrated circuit of claim 25, wherein the circuitry for identifying the image data region type includes, counting circuitry configured to compute a pixel value count within a neighborhood defined around a pixel of the image data; and

comparison circuitry configured to compare the pixel value count with a threshold value to determine a label associated with the image data region type.

29. The integrated circuit of claim 25, further comprising: descreening circuitry configured to blur halftone regions; and

enhancement circuitry configured to define edges over the blurred halftone regions.